

# Fair and Efficient Online Allocations\*

Gerdus Benadè

Boston University, Boston, MA 02215, benade@bu.edu

Aleksandr M. Kazachkov

University of Florida, Gainesville, FL 32611, akazachkov@ufl.edu

Ariel D. Procaccia

Harvard University, Boston, MA 02134, arielpro@seas.harvard.edu

Alexandros Psomas

Purdue University, West Lafayette, IN 47907, apsomas@cs.purdue.edu

David Zeng

Two Sigma, New York City, NY 10013, davidzengx@gmail.com

We study trade-offs between fairness and efficiency when allocating indivisible items online. We attempt to minimize envy, the extent to which any agent prefers another’s allocation to their own, while being Pareto efficient. We provide matching lower and upper bounds against a sequence of progressively weaker adversaries. Against worst-case adversaries we find a sharp trade-off: no allocation algorithm can simultaneously provide both non-trivial fairness and non-trivial efficiency guarantees. In a slightly weaker adversary regime where item values are drawn from (potentially correlated) distributions it is possible to achieve the best of both worlds. We give an algorithm that is Pareto efficient ex post and either envy-free up to one good or envy-free with high probability. Neither guarantee can be improved, even in isolation. En route, we give a constructive proof for a structural result of independent interest. Specifically, there always exists a Pareto efficient fractional allocation that is strongly envy-free with respect to pairs of agents with substantially different utilities, while allocating identical bundles to agents with identical utilities (up to multiplicative factors).

---

## 1. Introduction

Fairly and efficiently allocating resources to heterogeneous agents is a fundamental problem in operations research with applications including advertising ([Mehta et al. 2007](#), [Bateni et al. 2022](#), [Balseiro et al. 2021](#)), organ transplantation ([Su and Zenios 2006](#), [Bertsimas et al. 2013](#)), nurse shift

\* Preliminary versions of this paper were accepted at the 19th and 21st ACM Conference on Economics and Computation. Results from the first paper, on fair allocations in the worst-case, are included in Section 3. The second paper appeared as an extended abstract and considered fairness and efficiency against worst-case and Bayesian adversaries.

scheduling (Miller et al. 1976), and resource allocation in shared facilities like data centers (Butler and Williams 2002, Armony and Ward 2010, Ghodsi et al. 2011, Vardi et al. 2021).

We study the problem of allocating indivisible goods to agents who have additive valuations. Our goal is proving strong mathematical guarantees of both the *interpersonal fairness* and the *efficiency* of the resulting allocation. Several fairness notions have been used in the literature, but arguably the gold standard is *envy freeness*, which requires that each agent is at least as happy with their own allocation as the allocation of any other agent. In terms of efficiency, we aim for Pareto efficient or approximately Pareto efficient allocations which, in isolation, can be achieved by allocating each item to the agent who values it most.

Ignoring efficiency, envy-free solutions always exist in many well-studied fair division settings that involve *divisible* goods or a numéraire, such as cake cutting (Brams and Taylor 1996, Procaccia 2016) and rent division (Su 1999, Gal et al. 2017). For divisible items, one strategy for finding a fair allocation is the *competitive equilibrium from equal incomes* (CEEI) solution of Varian (1974). In the equilibrium allocation, agents use assigned (equal) budgets to purchase their preferred bundles of goods at virtual prices, and the market clears (all goods are allocated). This solution is envy free (Foley 1967) and coincides with the solution that maximizes the *Nash social welfare* (Arrow and Intriligator 1982), that is, the solution which maximizes the product of agent utilities.

By contrast, with *indivisible* goods, envy is clearly unavoidable in general — consider a single item that is desired by two agents. That is why previous papers (Lipton et al. 2004, Caragiannis et al. 2016) focus on the relaxed notion of *envy-freeness up to one good* (EF1), in which envy may exist, but for any bundle that an agent prefers over their own, there exists a single good whose removal eliminates that envy. With indivisible goods, the approximate-CEEI solution (Budish 2011) is EF1 but may not allocate all items, while the integral solution which maximizes the Nash social welfare is both EF1 and Pareto efficient (Caragiannis et al. 2016).

Our point of departure is that we allow items to arrive *online*. That is, we must choose how to allocate an item immediately and irrevocably at the moment it arrives, without knowing the values

of items that will arrive in the future. This setup mirrors common decision-making scenarios in humanitarian logistics. A paradigmatic example is that of food banks (Aleksandrov et al. 2015, Lee et al. 2019), which receive food donations and deliver them to nonprofit organizations such as food pantries and soup kitchens. Indeed, items are often perishable, which is why allocation decisions must be made quickly, and donated items are typically unsold or leftover products, leading to a lack of information about items that will arrive in the future.

As noted, the static setting permits a solution which is EF1 and Pareto efficient for any number of items (Caragiannis et al. 2016), but this requires upfront knowledge of all items. In contrast, in the online setting, one would expect the maximum envy to increase with the number of items but may hope to control the rate at which it grows. However, it is entirely unclear what impact minimizing envy online will have on efficiency. Our primary research question is:

*Are there online allocation algorithms that are simultaneously fair and efficient?*

### 1.1. Our Contributions

We study the tradeoff between fairness and efficiency in the following setting:  $T$  indivisible items arrive online (one-by-one) and must be allocated immediately and irrevocably to a set of agents  $\mathcal{N}$ . Agent  $i \in \mathcal{N}$  has value  $v_{it}$  for item  $t$ , these values are known at time of allocation and are generated according to one of four different adversary models, which we describe below. For each adversary model, we fully characterize the extent to which fairness and efficiency are compatible (or not).

In Section 3, we consider the strongest, worst-case adversaries. We start, in Section 3.1 and Section 3.2, by determining the limits of what is possible when solely minimizing envy with randomized allocation algorithms against an adaptive adversary that chooses the agent values for an arriving item after seeing the (realized) allocations of all the previous items. A natural idea is to allocate each item to an agent chosen uniformly at random. We find that this random allocation has *vanishing envy*: envy that grows sublinearly in the number of items (Theorem 1). Surprisingly, given the simplicity of the algorithm, we also construct a matching lower bound: Theorem 2 establishes

that the rate at which random allocation causes envy to vanish is asymptotically optimal (up to logarithmic factors). Unfortunately, random allocation only provides trivial efficiency guarantees.

Despite random allocation being asymptotically optimal in terms of fairness, there may exist other algorithms with vanishing envy that perform much better in terms of efficiency. We show that this is not the case. In Section 3.3, we study a weaker, *non-adaptive* worst-case adversary that selects an instance (with  $T$  items) after observing the algorithm, but before it is executed, so without knowledge of any random outcomes in the algorithm. Our main negative result (Theorem 3) is that, even against this weaker adversary, no algorithm with vanishing envy can have stronger efficiency guarantees than random allocation, implying the same result for adaptive adversaries. An important implication of Theorem 3 is that in settings where agents' value distributions are not known, or where there is a strong need for worst-case guarantees, algorithm designers are forced to choose between achieving either non-trivial efficiency guarantees or non-trivial fairness properties.

In Section 4, we study weaker, Bayesian adversaries. Section 4.1 considers the weakest of these, who selects a distribution  $D$  from which each value is drawn (independently and identically across items and agents). Here a good algorithm was identified by Dickerson et al. (2014) and later simplified and improved by Kurokawa et al. (2016), albeit in a different context: allocate each item to the agent who values it most. We find this core idea, with very minor modification, is ex post Pareto efficient and either envy-free with high probability or EF1 (Theorem 4).

When agents are non-identical, the strategy of allocating each item to the agent with the highest value fails, as do variants like considering the highest quantile instead of the highest value. Despite this, we design an algorithm that provides ex post Pareto efficiency and vanishing envy. Our main positive results are established against an even stronger adversary that allows for *correlated* agents, that is,  $v_{it}$  can be correlated with  $v_{i\hat{t}}$  but not with  $v_{i\bar{t}}$ . Of course, all results established for correlated agents extend to the settings with independent agents.

In Section 4.2, we analyze our high-level strategy, while postponing some crucial technical obstacles. We generate an offline instance with  $n$  agents and as many items as the support of the correlated discrete distribution  $D$ . We show in Theorem 5 that it is possible to use a (fractional) Pareto

efficient solution to this offline instance to guide the (integral) online allocation. This rounding can be coupled with any Pareto efficient and envy-free offline solution, for example the fractional allocation that maximizes the product of agents’ utilities, to yield an ex post Pareto efficient algorithm with vanishing envy.

Notably, if the solution to the offline instance is a *strongly envy-free* allocation, where each agent strictly prefers their own allocation over any other, the same approach would imply online envy-freeness with high probability (a much stronger guarantee than simply vanishing envy). This goal is too optimistic. However, we show in Section 4.3 that it is possible to provide an offline allocation with a slightly weaker property, which, when used online, results in either envy-freeness with high probability or EF1 ex post (Theorem 8). Remarkably, this is the same guarantee as against the weak Bayesian adversary.

Theorem 8 relies on a structural, constructive result about fractional allocations to the offline problem (Theorem 7). We give an algorithm that starts with a solution to the Eisenberg-Gale convex program (Eisenberg and Gale 1959) with equal budgets and iteratively adjusts the budgets until it arrives at a Pareto efficient fractional allocation where agent  $i$  either strictly prefers their allocation to the allocation of agent  $j$ , or, if they are indifferent, then  $i$  and  $j$  have *identical* fractional allocations and the *same* value (up to multiplicative factors) for all items allocated to them. We believe this result and approach may be of independent interest.

We conclude with a remark on the fairness criteria of our main positive result: “EF1 or envy-free with high probability”. Even in isolation under the weakest adversary, this is the strongest achievable fairness guarantee. It is impossible to *always* output an EF1 allocation (ex post), and it is impossible to *always* output an allocation that is envy free with high probability (see Section EC.2).

## 1.2. Related Work

Our paper is related to the growing literature on *online* or *dynamic* fair division (He et al. (2019), Kash et al. (2014), Friedman et al. (2015, 2017), Li et al. (2018), Freeman et al. (2018), Aleksandrov et al. (2015), Walsh (2011), Bogomolnaia et al. (2021), Gkatzelis et al. (2021)). In settings similar

to our worst-case adversary, He et al. (2019) allow items to be reallocated at a later time, and study the number of *adjustments* that are necessary and sufficient in order to maintain an EF1 allocation online. Bansal et al. (2020) propose an algorithm that guarantees envy of  $O(\log T)$  with high probability for the case of two independent identical agents but do not consider efficiency. In contrast to our positive result in Section 4.3, their result allows the distribution to depend on  $T$ .

Dickerson et al. (2014) study a completely different setting and show that allocating an item to the agent who values it most results in an envy-free allocation with probability 1 as the number of items goes to infinity (a similar result appears in Kurokawa et al. (2016)). It is straightforward to apply this against the weakest adversary we consider, where agents are identical and items values are independent and identically distributed. We discuss their result in greater detail in Section 4.

For the offline problem, i.e., when all agents' values are available to the algorithm, Caragiannis et al. (2016) show that in fact there is no tradeoff between fairness and efficiency: the (integral) allocation that maximizes the Nash social welfare is simultaneously Pareto efficient and EF1. Computing the fractional allocation that maximizes Nash social welfare is a special case of the Fisher market equilibrium with affine utility buyers; the latter problem was solved in (weakly) polynomial time by Devanur et al. (2008) and improved to a strongly polynomial time algorithm by Orlin (2010). Our structural result starts from an exact solution to the Eisenberg-Gale convex program (Eisenberg and Gale 1959) and then uses a polynomial number of operations. Therefore, all our algorithms run in strongly polynomial time; we further comment on this in Section 5.

Beyond envy, the *price of fairness* measures the relative loss in social welfare that result from enforcing a fairness constraint. The price of fairness has been studied in static settings for divisible (Caragiannis et al. 2009, Bertsimas et al. 2011, 2012) and more recently indivisible items (Barman et al. 2020, Bei et al. 2021, Narayan et al. 2021). This is similar in spirit to our work; we approximate Pareto efficiency rather than welfare and are willing to relax the fairness notion rather than strictly enforcing it.

## 2. Preliminaries

We study the problem of allocating a set of  $T$  indivisible items (also referred to as goods) arriving over time, labeled by  $\mathcal{G} = [T] = \{1, 2, \dots, T\}$ , to a set of  $n$  agents, labeled  $\mathcal{N} = [n]$ . Agent  $i \in \mathcal{N}$  assigns a (normalized) value  $v_{it} \in [0, 1]$  to each item  $t \in \mathcal{G}$ . Agents have additive utilities for subsets of items, where  $v_i(S) = \sum_{t \in S} v_{it}$  for  $S \subseteq \mathcal{G}$ . An *allocation*  $A$  is a partition of the items into bundles  $A_1, \dots, A_n$ , where  $A_i$  is assigned to agent  $i \in \mathcal{N}$ .

Items arrive one by one, in order, over a total of  $T$  rounds and are immediately allocated. Let  $\mathcal{G}^t = [t]$  be the set of items that have arrived up until time  $t$ . Allocations of  $\mathcal{G}^t$  are denoted  $A^t$ . Agents' valuations for the  $t$ -th item only become available once the item arrives, and we would like to allocate the goods so that the final allocation  $A = A^T$  is *fair* and *efficient*. Many of our results characterize fairness and efficiency as  $T$  grows. We use standard asymptotic notation; see Appendix [EC.1](#) for a reminder.

We now discuss the different adversary models which govern how the item values are generated before formally defining our notions of fairness and efficiency.

### 2.1. Adversary Models

One may think of each scenario as a game between the adversary and the allocation algorithm. For the first two, it will be convenient to think of the algorithm being fixed before the adversary picks a strategy. For the last two adversaries, it will be more intuitive to think of the adversary picking a strategy (distribution) first.

We list our adversaries from strongest to weakest, where a stronger adversary can simulate the strategy of a weaker adversary but not vice versa. Distributions are assumed to be discrete with finite support and independent of  $T$ , so it cannot have support of size  $T$ , variance  $1/T$ , etc. We refer to adversaries (1)–(2) as *worst-case*, and (3)–(4) as *Bayesian*.

1. **Adaptive Adversary.** The adversary selects values  $\{v_{it}\}_{i \in \mathcal{N}}$  after observing the algorithm's allocations for the first  $t - 1$  items.

2. **Non-adaptive adversary.** The adversary selects an instance (with  $n$  agents and  $T$  items) after seeing the algorithm's description, but without knowing the outcome of any randomness in the algorithm. Our main negative result is for this setting.
3. **Correlated agents and i.i.d. items.** The adversary specifies a joint distribution for agent values  $D_1, \dots, D_n$ . In round  $t$ , the value of item  $t$  to each agent  $i$  is drawn from their distribution, that is  $v_{it} \sim D_i$ . Value  $v_{it}$  can be correlated with  $v_{jt}$ , but not with  $v_{i\hat{t}}$ . For simplicity, we treat this setting as follows. Each item  $t$  has one of  $m$  types. Agent  $i$  has value  $v_i(\gamma_j)$  for an item of type  $\gamma_j$ ; the type of each item is drawn i.i.d. from a distribution  $D$  with support  $G_D$ ,  $|G_D| = m$ . We write  $f_D(\gamma_j)$  for the probability that the  $t$ -th item has type  $\gamma_j$ . Our main positive result is for this setting.
4. **Identical agents and i.i.d. items.** The adversary selects a distribution  $D$ . In round  $t$ , the value of item  $t$  to each agent  $i$  is drawn independently from this distribution, i.e.  $v_{it} \sim D$ .

Against Bayesian adversaries we study the allocation algorithm's performance as  $T \rightarrow \infty$ . Worst-case adversaries always have the option to let all future items be worthless to every agent, so here  $T$  is assumed to be fixed and known when the adversary selects their strategy.

## 2.2. Measuring Fairness

We focus on a well-studied notion of fairness called *envy*. An allocation  $A = (A_1, \dots, A_n)$  is *envy-free* when  $v_i(A_i) \geq v_i(A_j)$  for all  $i, j \in \mathcal{N}$ . The pairwise envy of agent  $i$  towards  $j$  is  $\text{Envy}_{i,j}(A) = \max\{v_i(A_j) - v_i(A_i), 0\}$ , while  $\text{Envy}(A) = \max_{i,j \in \mathcal{N}} \text{Envy}_{i,j}(A)$  is the maximum envy.  $\text{Envy}(A) = 0$  implies the allocation is envy free. An allocation  $A$  is *envy-free up to one good* (EF1) when, for all pairs of agents  $i, j$ ,  $\text{Envy}_{i,j}(A) \leq \max_{t \in A_j} v_{it}$ . Note that this is a stronger guarantee than  $\text{Envy}(A) \leq 1$  when  $\max_{t \in \mathcal{G}} v_{it} < 1$ . For convenience, we will occasionally refer to  $\text{Envy}(A^k)$  as  $\text{Envy}_k$  for  $k \in \mathcal{G}$ , and  $\text{Envy}_T = \text{Envy}(A^T) = \text{Envy}(A)$ . An algorithm has *vanishing envy* if the expected maximum pairwise envy is sublinear in  $T$ , that is,  $\mathbb{E}[\text{Envy}(A)] \in o(T)$  or  $\lim_{T \rightarrow \infty} \mathbb{E}[\text{Envy}(A)]/T \rightarrow 0$ .



### 2.3. Measuring Efficiency

The *utility profile* of an allocation  $A$  is a vector  $u = (u_1, \dots, u_n)$  where  $u_i = v_i(A_i)$ . A utility vector  $u$  *dominates* another utility vector  $u'$ , denoted by  $u \succ u'$ , if  $u_i \geq u'_i$  for all  $i$  and there is some  $j$  for which  $u_j > u'_j$ . An allocation with utility profile  $u$  is *Pareto efficient* if there is no allocation with utility vector  $u'$  such that  $u' \succ u$ . Where appropriate, we use a notion of approximate Pareto efficiency, initially defined by [Ruhe and Fruhwirth \(1990\)](#), to measure the efficiency of our algorithms. An allocation with utility profile  $u$  is  $\alpha$ -Pareto efficient when  $u/\alpha$  is undominated.

Since our setting is online, we need to specify whether efficiency guarantees are worst-case or average-case with respect to the adversary instance and the randomness of our algorithms. For a worst-case guarantee, we say that an allocation is  $\alpha$ -Pareto efficient *ex post* if it always outputs an  $\alpha$ -Pareto efficient allocation, that is, for all agent valuations and all possible outcomes of any randomness in the algorithm. On the other hand, an allocation algorithm is  $\alpha$ -Pareto efficient *ex ante* if the expected utility profile is  $\alpha$ -Pareto efficient (where the expectation is with respect to the randomness in the instance and the algorithm). Our main positive result guarantees 1-Pareto efficiency *ex post*, while our main negative result shows that a specific notion of fairness is incompatible with  $1/n$ -Pareto efficiency *ex ante*.

### 3. Fairness and Efficiency are Incompatible Against Worst-case Adversaries

In this section we discuss the trade-off between efficiency and fairness against the stronger, non-Bayesian adversaries.

To build intuition, we consider a couple of obvious strategies for finding fair or efficient allocation algorithms and highlight how they fail. First, we observe that the natural Pareto efficient algorithm that allocates each item to the agent who values it most has  $\text{Envy}_T \in \Omega(T)$ .

**EXAMPLE 1.** Consider two agents. Let  $v_{1t} = 1$  for all  $t \in \mathcal{G}$ , and  $v_{2t} = 1/2$  for all  $t \in \mathcal{G}$ . When allocating each item to the agent who likes it most,  $A_1 = \mathcal{G}$  and  $A_2 = \emptyset$ . This allocation is Pareto efficient, but has  $\text{Envy}_T = T/2$ .

**Table 1** Blindly allocating to the agent with the highest envy leads to constant per-round envy.

$t$	1	2	3	4	5	...
Value of agent 1	$\boxed{1/2}$	1	$\boxed{\epsilon}$	1	$\boxed{\epsilon}$	...
Value of agent 2	1/2	$\boxed{\epsilon}$	1	$\boxed{\epsilon}$	1	...
Envy of agent 1	-1/2	1/2	$1/2 - \epsilon$	$3/2 - \epsilon$	$3/2 - 2\epsilon$	...
Envy of agent 2	1/2	$1/2 - \epsilon$	$3/2 - \epsilon$	$3/2 - 2\epsilon$	$5/2 - 2\epsilon$	...

The prior allocation algorithm ignored envy entirely, so it is no surprise that it had linear envy. Our next example analyzes a greedy policy that allocates each item to the agent with the greatest envy and finds it, too, fails to achieve vanishing envy.

**EXAMPLE 2.** Consider the algorithm that at step  $t$  allocates the item to the agent with the maximum envy, if she has positive value for the item, and otherwise, say, allocates to the agent with the highest value for the item. We claim this algorithm can lead to  $\text{Envy}_T \in \Omega(T)$ .

We construct an example where each agent envies the other after the second item is allocated. For  $t \geq 2$ , whenever agent  $i$  has maximum envy, we present an item with value  $\epsilon$  for her, and value 1 for the other agent. Table 1 summarizes the analysis.

For  $t \geq 2$ , the envy of each agent increases by 1 every two steps. Therefore, the maximum envy at step  $2t$  is approximately  $t$ , and  $\text{Envy}_T/T$  approaches  $1/2$  as  $T$  goes to infinity.

These examples suggest it is non-trivial to come up with an allocation algorithm that achieves vanishing envy. Couple vanishing envy with Pareto efficiency, and the task appears quite daunting.

We first investigate what is possible when focusing solely on fairness. We find that vanishing envy is achievable; in fact, uniform random allocation has  $\mathbb{E}[\text{Envy}_T] \in \tilde{O}(\sqrt{T/n})$  against adaptive adversaries, while trivially being  $\frac{1}{n}$ -Pareto efficient ex ante. Now the question becomes: is this optimal, or are there other strategies with even stronger fairness properties? We provide an adaptive adversary strategy which guarantees  $\text{Envy}_T \in \Omega((T/n)^{r/2})$  for any  $r < 1$ , thereby showing that random allocation is optimal (up to logarithmic factors) in terms of envy.

Finally, we turn our attention to simultaneously providing fairness and efficiency guarantees. We find that, even against a non-adaptive adversary, no algorithm can achieve vanishing envy

while being  $(\frac{1}{n} + \varepsilon)$ -Pareto efficient for any  $\varepsilon > 0$ . This clearly establishes the boundaries of what is possible against worst-case adversaries: any allocation algorithm must choose between achieving *either* non-trivial fairness guarantees *or* non-trivial efficiency.

### 3.1. Random Allocation has Vanishing Envy and is $1/n$ -Pareto Efficient

A natural randomized algorithm is to allocate each item (independently) to an agent selected uniformly at random; we refer to this as the *random allocation algorithm*. The following observation is a direct result of the fact that each agent receives each item with probability  $1/n$  under random allocation and therefore has expected utility  $1/n$  times their utility for all items.

PROPOSITION 1. *The random allocation algorithm is  $1/n$ -Pareto efficient ex ante.*

Next, we analyze the fairness of the random allocation algorithm by first characterizing the adversary's optimal strategy. We prove that for an adaptive adversary who maximizes  $\mathbb{E}[\text{Envy}_T]$ , where the expectation is with respect to the randomness of the algorithm, the optimal strategy is integral, that is, all values are in  $\{0, 1\}$ . In fact, the optimal integral strategy sets assigns  $v_{it} = 1$  for all  $i \in [n], t \in [T]$ . This optimal adversary strategy is non-adaptive and therefore, since all the randomness is coming from the algorithm, the random variables for the envy between agents  $i$  and  $j$  at times  $t$  and  $t'$  are independent. Standard concentration inequalities for the envy between any pair of agents, combined with a union bound over all such pairs, gives an upper bound on the expected envy.

THEOREM 1. *Suppose that  $T \geq n \log T$ , where  $\log$  is the natural logarithm. Then the random allocation algorithm guarantees that  $\mathbb{E}[\text{Envy}_T] \in O(\sqrt{T \log T / n})$ .*

The assumption of  $T \geq n \log T$  is innocuous, as otherwise we can give each agent at most  $\log T$  items to achieve  $\text{Envy}_T \leq \log T$ .

Proof of Theorem 1. A typical extensive-form game tree would have nodes associated with the algorithm or the adversary, and arcs corresponding to actions (the allocation of the current item in the case of the algorithm, choosing a value vector in the case of the adversary). However,

because we consider a fixed algorithm, it is convenient to imagine an unusual, adversary-oriented game tree.

Consider a game tree with nodes on  $T + 1$  levels. Every node on level  $1, \dots, T$  has  $n$  outgoing arcs labeled  $1, \dots, n$ . The leaf nodes on level  $T + 1$  are labeled by the maximum envy for the corresponding path, which defines an allocation of the  $T$  items.

A fully adaptive strategy  $s$  for the adversary is defined by labeling every internal node  $u$  with a value vector  $s(u)$ , where  $s(u)_i$  is the value of agent  $i$  for the item corresponding to node  $u$ . The adversary's strategy is allowed to depend on the allocations and valuations so far, i.e., the path from the root to  $u$ . The objective of the adversary is to choose a strategy  $s$  that maximizes the expected envy. The algorithm selects an outgoing edge at every node  $u$ , corresponding to an allocation of the item with valuation  $s(u)$ . Consider the algorithm that allocates every item uniformly at random or, equivalently, picks a random outgoing edge at each node  $u$ .

The following two lemmas are inspired by the work of Sanders (1996) on load balancing and show that the adversary labels every internal node of this tree with the vector  $\mathbf{1}^n$ . All omitted proofs appear in the electronic companion.

LEMMA 1. *For every allocation algorithm, the adversary has an optimal adaptive strategy that labels every internal node of the game tree with a vector in  $\{0, 1\}^n$ .*

This holds for any allocation algorithm, since for every agent's valuation of any item it is possible to compute whether that item increases or decreases the maximum envy in expectation. If it increases (resp. decreases) the maximum envy, the adversary benefits by increasing (resp. decreasing) the corresponding valuation to 1 (resp. to 0).

The following lemma leverages specific properties of the random allocation algorithm.

LEMMA 2. *Against uniformly random allocations, the adversary has an optimal adaptive strategy that labels every internal node of the game tree with the vector  $\mathbf{1}^n$ .*

The fact that the adversary is adaptive naturally introduces a dependence in the change in any pairwise envy from one arrival to the next. Lemma 2 allows us to circumvent this dependence as

though we are dealing with a non-adaptive adversary and express any pairwise envy as the sum of independent random variables.

Specifically, given this adversary strategy, define independent random variables

$$X_t^{ij} = \begin{cases} -1, & \text{with probability } 1/n, \\ 0, & \text{with probability } 1 - 2/n, \\ 1, & \text{with probability } 1/n \end{cases}$$

for all  $t \in [T]$ ,  $i, j \in [n]$ . Clearly,  $\text{Envy}_T^{ij} = \max_{i, j \in [n]} \{\sum_{t=1}^T X_t^{ij}, 0\}$ . For each  $X_t^{ij}$ ,  $\mathbb{E}[X_t^{ij}] = 0$ ,  $\mathbb{E}[(X_t^{ij})^2] = 2/n$  and  $|X_t^{ij}| \leq 1$ . We bound the probability of having large envy between any pair of agents  $i$  and  $j$  by applying Bernstein's inequality ((Bernstein 1946), see EC.3) to  $\text{Envy}_T^{ij}$ , which equals  $\sum_{t=1}^T X_t^{ij}$  when envy exists. It follows that, for  $\lambda > 0$ ,

$$\Pr[\text{Envy}_T^{ij} \geq \lambda] = \Pr\left[\sum_{t=1}^T X_t^{ij} \geq \lambda\right] \leq \exp\left(-\frac{\frac{1}{2}\lambda^2}{\frac{2T}{n} + \frac{1}{3}\lambda}\right) = \exp\left(-\frac{3n\lambda^2}{12T + 2\lambda n}\right).$$

Let  $\lambda = 10\sqrt{T \log T/n}$ . Taking a union bound over pairs of agents gives

$$\Pr[\text{Envy}_T \geq \lambda] = \Pr[\exists i, j \in [n] \text{ such that } \text{Envy}_T^{ij} \geq \lambda] \leq n^2 \exp\left(-\frac{300T \log T}{12T + 20\sqrt{nT \log T}}\right) \leq \frac{1}{T},$$

where the last inequality uses the assumption that  $T \geq n \log T$ . Since the maximum possible envy is  $T$ , the desired bound on expected envy directly follows, completing the proof of Theorem 1.  $\square$

The existence of a randomized algorithm with  $\text{Envy}_T \in O(\sqrt{T \log T/n})$  implies the existence of deterministic algorithms with the same guarantee. One such algorithm can be found through standard derandomization techniques (Alon and Spencer 2000). This deterministic algorithm can be interpreted as placing an exponential penalty on each pairwise envy and greedily allocating each item to minimize the sum of penalties at the end of each round (Benadè et al. 2018).

### 3.2. Random Allocation Optimizes Fairness Against Adaptive Adversaries

In this section, we show that an adversary can guarantee  $\text{Envy}_T \in \Omega((T/n)^{r/2})$  for any  $r < 1$ . As  $r \rightarrow 1$ , it follows that the random allocation algorithm in Section 3.1 is optimal (up to a logarithmic factor).

**THEOREM 2.** *For any  $n \geq 2$  and  $r < 1$ , there exists an adversary strategy for setting item values such that any algorithm must have  $\text{Envy}_T \in \Omega((T/n)^{r/2})$ .*


We first prove the bound for  $n = 2$ , followed by the case of an arbitrary number of agents.

**LEMMA 3.** *For  $n = 2$  and any  $r < 1$ , there exists an adversary strategy for setting item values such that any algorithm must have  $\text{Envy}_T \in \Omega(T^{r/2})$ .*

*Proof.* Label the agents  $L$  and  $R$ , and let  $\{v_0 = 1, v_1, v_2, \dots\}$  be a decreasing sequence of values (specified later) satisfying  $v_d - v_{d+1} < v_{d'} - v_{d'+1}$  for all  $d' < d$ . The adversary keeps track of the *state* of the game, and the current state defines its strategy for choosing the agents' valuations for the next item. The lower bound follows from the adversary strategy illustrated in Figure 1. Start in state 0, which we will also refer to as  $L_0$  and  $R_0$ , where the adversary sets the value of the arriving item as  $(1, 1)$ . To the left of state 0 are states labeled  $L_1, L_2, \dots$ ; when in state  $L_d$ , the next item that arrives has value  $(1, v_d)$ . To the right of state 0 are states labeled  $R_1, R_2, \dots$ ; when in state  $R_d$  the next item arrives with value  $(v_d, 1)$ . Whenever the algorithm allocates an item to agent  $L$  (resp.  $R$ ), which we will refer to as making an  $L$  (resp.  $R$ ) step, the adversary moves one state to the left (resp. right).

We construct the optimal allocation algorithm against this adversary, and show that for this algorithm the envy at some time step  $t \in [T]$  will be at least  $\Omega(T^{r/2})$  for the given  $r < 1$ . This immediately implies Lemma 3: if the envy is sufficiently large at any time step  $t$  the adversary can guarantee the same envy at time  $T$  by making all future items valued at zero by both agents.

**Figure 1** Adversary strategy for two-agent lower bound. In state  $L_d$ , an item valued  $(1, v_d)$  arrives, while in state  $R_d$ , an item valued  $(v_d, 1)$  arrives. The arrows indicate whether agent  $L$  or agent  $R$  is given the item in each state. The arrows are labeled by the amount envy changes after that item is allocated.



The intuition for the adversary strategy we have defined is that it forces the algorithm to avoid entering state  $L_d$  or  $R_d$  for high  $d$ , as otherwise the envy of some agent will grow to  $v_0 + v_1 + \dots + v_d$ , which will be large by our choice of  $\{v_d\}$ . At the same time, if an  $L$  step is taken at state  $L_d$ , followed by a later return to state  $L_d$ , the envy of  $R$  increases by at least  $v_d - v_{d+1}$ ; we choose  $\{v_d\}$  so that this increase in envy is large enough to ensure that any algorithm which spends too many time steps close to state 0 incurs large envy.

By the pigeonhole principle, either the states to the left or to the right of state 0 are visited for at least half the time. Assume, without loss of generality, that our optimal algorithm spends time  $T' = \lceil T/2 \rceil$  in the ‘left’ states  $(L_0, L_1, \dots)$ , and that  $T'$  is even. We prove that the envy of agent  $R$  grows large at some time step  $t$ . We ignore any time the algorithm spends in the states  $R_d$ ,  $d \geq 1$ . To see why this is without loss of generality, consider first a cycle spent in the right states that starts at  $R_0$  with an item allocated to  $R$  and eventually returns to  $R_0$ . In such a cycle, an equal number of items are allocated to both agents. All of these items have value 1 to agent  $R$ , yielding a net effect of 0 on agent  $R$ ’s envy. (We ignore agent  $L$  completely, as our analysis is of the envy of agent  $R$ .) The other case is when the algorithm starts at  $R_0$  but does not return to  $R_0$ . This scenario can only occur once, which means that the algorithm has already taken  $T'$  steps on the left side; the allocation of these items does not affect our proof.

Let  $0 \leq K \leq T'/2$  be an integer and denote by  $\text{OPT}(K)$  the set of envy-minimizing allocation algorithms that spend the  $T'$  steps in states  $L_0, \dots, L_K$  (and reach  $L_K$ ). Note that the algorithm aims to minimize the maximum envy at any point in its execution. Let  $\mathcal{A}^*(K)$  be the following algorithm, starting at  $L_0$ : Allocate the first  $K$  items to agent  $L$ , thus arriving at state  $L_K$ . Alternate between allocating to agents  $R$  and  $L$  for the next  $T' - 2K$  items, thereby alternating between states  $L_{K-1}$  and  $L_K$ . Allocate the remaining  $K$  items to agent  $R$ . Our first result is that  $\mathcal{A}^*(K)$  belongs to  $\text{OPT}(K)$ .

LEMMA 4.  $\mathcal{A}^*(K) \in \text{OPT}(K)$ .

We analyze the envy of  $\mathcal{A}^*(K)$  as a function of  $K$  before optimizing  $K$ . Agent  $R$ ’s maximum envy is realized at step  $T' - K$ , right before the sequence of  $R$  moves.  $\text{Envy}_{T'-K}$  has two terms:

the envy accumulated to reach state  $L_K$ , and the envy from alternating  $R$  and  $L$  moves between states  $L_K$  and  $L_{K-1}$ , so

$$\text{Envy}_{T'-K} = \sum_{d=0}^{K-1} v_d + \frac{T' - 2K}{2} \cdot (v_{K-1} - v_K). \quad (1)$$

Given  $r < 1$ , define  $v_d = (d+1)^r - d^r$ . Notice that  $\sum_{d=0}^{K-1} v_d = K^r$ . When  $K \geq \sqrt{T'/2}$  it follows that  $\sum_{d=0}^{K-1} v_d \geq (T'/2)^{r/2} \in \Omega(T'^{r/2})$ , which is what we set out to prove. We limit the rest of the analysis to the case where  $K \leq \sqrt{T'/2}$ .

LEMMA 5. *Let  $K \leq \sqrt{T'/2}$  and define  $v_d = (d+1)^r - d^r$  for  $r < 1$ . Then  $v_{K-1} - v_K \geq r(1-r)K^{r-2}$ .*

Applying Lemma 5 to (1) and distributing terms yields

$$\text{Envy}_{T'-K} \geq K^r - r(1-r)K^{r-1} + \frac{T'}{2}r(1-r)K^{r-2} \geq \frac{1}{2}(K^r + T'r(1-r)K^{r-2}), \quad (2)$$

where the second inequality uses the fact that  $r(1-r) \leq 1/4 < 1/2$  and assumes  $K > 1$  (otherwise the envy would be linear in  $T'$ ). To optimize  $K$ , noting that the second derivative of the above bound is positive for  $K \leq \sqrt{T'/2}$ , we find the critical point:

$$\frac{\partial}{\partial K} (K^r + T'r(1-r)K^{r-2}) = rK^{r-1} - T'r(1-r)(2-r)K^{r-3} = 0 \implies K = \sqrt{T'(1-r)(2-r)}.$$

Defining  $C_1 = \sqrt{(1-r)(2-r)}$  and substitute into (2) to obtain

$$\text{Envy}_{T'-K} \geq \frac{1}{2} (C_1^r (T')^{r/2} + T'r(1-r)C_1^{r-2} (T')^{r/2-1}) \in \Omega(T'^{r/2}). \quad \square$$

The extension to  $n$  agents follows from the same strategy for agents  $L$ ,  $R$  and letting all other agents value every item at zero. Allocating to an agent  $i \in \mathcal{N} \setminus \{L, R\}$  does not advance the state.

### 3.3. Non-Trivial Fairness and Efficiency are Incompatible

Recall that random allocation was  $\frac{1}{n}$ -Pareto efficient. We conclude this section by showing that no algorithm with vanishing envy can improve on this efficiency guarantee against a non-adaptive worst-case adversary, which immediately establishes the result against adaptive adversaries.



**THEOREM 3.** *Against a non-adaptive adversary, no (randomized or deterministic) allocation algorithm can achieve both  $\text{Envy}_T \in o(T)$  and be  $(\frac{1}{n} + \varepsilon)$ -Pareto efficient ex ante, for any  $\varepsilon > 0$ .*

To build up some intuition, we start by considering the case of an adaptive adversary where the algorithm must achieve vanishing envy and  $(\frac{1}{n} + \varepsilon)$ -Pareto efficiency ex post. Recall that randomization does not help against an adaptive adversary, so we focus on deterministic algorithms.

**LEMMA 6.** *No deterministic allocation algorithm can achieve both  $\text{Envy}_T \in o(T)$  and be  $(\frac{1}{n} + \varepsilon)$ -Pareto efficient ex post, for any  $\varepsilon > 0$ , against an adaptive adversary.*

*Proof.* Consider any vanishing envy algorithm that for any given  $T$  produces an allocation  $A^T$ , where  $\text{Envy}(A^T) \leq f(T)$  for some  $f(T) \in o(T)$ , and assume, for the sake of contradiction, that this algorithm achieves  $(\frac{1}{n} + \varepsilon)$ -Pareto efficiency for some  $\varepsilon > 0$ .

We construct an instance denoted  $I$ , parameterized by  $\varepsilon$  and  $T$ , which will lead to a contradiction. For each agent  $i \in \mathcal{N}$ ,  $v_{ij} = 1$  for  $j \in [\frac{T}{n}(i-1) + 1, \dots, \frac{T}{n}i]$  and all other items  $j'$  have value  $v_{ij'} = \varepsilon$ . Notice that agent  $i$  cares chiefly about the  $i$ -th segment of  $T/n$  items.

Note that for all intermediate allocations at time  $t \leq T$ , we must still have  $\text{Envy}(A^t) \leq f(T)$  since an adaptive adversary could always make the remaining items valueless to all agents. The first step is to show via induction that for all “segments” of items  $[\frac{T}{n}(i-1) + 1, \dots, \frac{T}{n}i]$ , every agent must receive a number of items in  $[\frac{T}{n^2} - x_i, \frac{T}{n^2} + x_i]$ , where  $x_i = \frac{f(T)}{\varepsilon} (1 + \frac{2}{\varepsilon})^{i-1}$ .

As base case for the inductive argument, consider the first segment (*i.e.*  $i = 1$ ). Suppose that some agent  $k$  receives  $\frac{T}{n^2} + y$  items where  $y > 0$ . Another agent  $\hat{k}$  must then receive fewer than  $\frac{T}{n^2}$  items. Then, the envy of  $\hat{k}$  for  $k$  at the end of the first segment,  $\text{Envy}_{\hat{k},k}(A^{T/n})$  is at least  $\varepsilon \cdot y$ . But,  $\text{Envy}_{\hat{k},k}(A^{T/n}) \leq f(T)$ , which implies that  $y \leq \frac{f(T)}{\varepsilon}$ ; the lower bound on  $y$  is identical.

For the inductive step, again suppose that in the segment  $[\frac{T}{n}(i-1) + 1, \dots, \frac{T}{n}i]$  some agent  $k$  receives  $\frac{T}{n^2} + y$  items, where  $y > 0$ , and let  $\hat{k}$  be the agent who received fewer than  $\frac{T}{n^2}$  items. At the start of segment  $i$ ,

$$v_{\hat{k}} \left( A_k^{\frac{T}{n}(i-1)} \right) - v_{\hat{k}} \left( A_{\hat{k}}^{\frac{T}{n}(i-1)} \right) \geq - \sum_{i' < i} 2x_{i'},$$

which occurs when  $\hat{k}$  received  $\frac{T}{n^2} + x_{i'}$  items from each segment,  $k$  got  $\frac{T}{n^2} - x_{i'}$ , and  $\hat{k}$  had value 1 for all items up until  $\frac{T}{n}(i-1)$ . Therefore, after the  $i$ -th segment,

$$f(T) \geq \text{Envy}_{\hat{k},k}(A^{\frac{T}{n}i}) \geq \varepsilon \cdot y + v_{\hat{k}}\left(A_k^{\frac{T}{n}(i-1)}\right) - v_k\left(A_{\hat{k}}^{\frac{T}{n}(i-1)}\right) \geq \varepsilon \cdot y - 2 \sum_{i' < i} x_{i'},$$

which implies that

$$\begin{aligned} y &\leq \frac{1}{\varepsilon} \left( f(T) + 2 \sum_{i' < i} x_{i'} \right) = \frac{1}{\varepsilon} \left( f(T) + 2 \sum_{i' < i} \frac{f(T)}{\varepsilon} \left( 1 + \frac{2}{\varepsilon} \right)^{i'-1} \right) \\ &= \frac{f(T)}{\varepsilon} \left( 1 + 2 \sum_{p=0}^{i-2} \left( 1 + \frac{2}{\varepsilon} \right)^p \right) = \frac{f(T)}{\varepsilon} \left( 1 + \frac{2}{\varepsilon} \right)^{i-1}, \end{aligned}$$

where the final transition results from summing the geometric series. The bound on  $y$  is identical when we consider the case that  $y < 0$ .

Next, we show that the allocation  $A^T$  cannot be  $(\frac{1}{n} + \varepsilon)$ -Pareto efficient. First, note that the social welfare maximizing allocation achieves utility  $(\frac{T}{n}, \dots, \frac{T}{n})$  by giving all the items of the  $i$ -th segment to agent  $i$ . Meanwhile, since  $x_i < x_n$ , we have that in  $A^T$  each agent gets utility  $u_i$  at most  $(1 + (n-1)\varepsilon)(\frac{T}{n^2} + x_n)$ . Therefore,

$$\begin{aligned} \frac{u_i}{1/n + \varepsilon} &< (1 + (n-1)\varepsilon) \left( \frac{T}{n^2} + x_n \right) \left( \frac{1}{\frac{1}{n} + \varepsilon} \right) = (1 + (n-1)\varepsilon) \left( \frac{T}{n^2} + \frac{f(T)}{\varepsilon} \left( 1 + \frac{2}{\varepsilon} \right)^{n-1} \right) \frac{n}{1 + \varepsilon n} \\ &= \frac{1 + (n-1)\varepsilon}{1 + \varepsilon n} \cdot \left( \frac{T}{n} + n \cdot \frac{f(T)}{\varepsilon} \left( 1 + \frac{2}{\varepsilon} \right)^{n-1} \right) \\ &= \frac{T}{n} \cdot \left( 1 - \frac{\varepsilon}{1 + \varepsilon n} \right) \cdot \left( 1 + \frac{f(T)}{T} \cdot \frac{n^2}{\varepsilon} \left( 1 + \frac{2}{\varepsilon} \right)^{n-1} \right). \end{aligned}$$

For large enough  $T$ , in particular when  $\frac{f(T)}{T} < \frac{\varepsilon}{1 + (n-1)\varepsilon} \cdot \frac{\varepsilon}{n^2(1+2/\varepsilon)^{n-1}}$ , this implies  $u_i < \frac{T}{n} \cdot (1/n + \varepsilon)$  for each agent  $i$ . We conclude  $A^T$  is not  $(\frac{1}{n} + \varepsilon)$ -Pareto efficient, a contradiction.  $\square$

We use this result to prove Theorem 3 for a non-adaptive adversary.

Proof of Theorem 3. Suppose that there is an allocation algorithm which guarantees that for any  $T$ , no matter the instance the adversary selects,  $\mathbb{E}[\text{Envy}(A^T)] \leq f(T)$  for some  $f(T) \in o(T)$ , where the expectation is over the randomness used by the algorithm. We will describe a family of  $n$  instances. After the arrival of the first  $\frac{T}{n}i$  items, it will be impossible for the allocation algorithm

to distinguish between  $n - i + 1$  of these instances. For  $i \in \{1, \dots, n\}$ , instance  $I_i$ 's first  $\frac{T}{n}i$  items follow  $I$ , the instance of the adaptive adversary described above, and the remaining items have no value. Let  $g(T)$  be a function such that  $g(T) \cdot f(T) \in o(T)$  and  $g(T) \in \omega(1)$ . For example, one may think of  $g(T) = T^\delta$ , for some small  $\delta > 0$  that depends on  $f(T)$ .

Again, we bound the number of items the algorithm can allocate to each agent in each segment; this time our bounds will be looser and probabilistic. Consider the behavior of the algorithm when faced with instance  $I_1$ . At the end of the first segment, i.e. for items 1 through  $\frac{T}{n}$ , if the algorithm allocates to some agent  $k$  at least  $\frac{T}{n^2} + x_1$ , for  $x_1 > 0$ , with probability at least  $\frac{1}{g(T)}$ , then the expected envy of some agent  $\hat{k}$  (the one who received fewer than  $\frac{T}{n^2}$  items) is at least  $\varepsilon \cdot x_1$ , that is  $\mathbb{E}[\text{Envy}(A^T)] \geq \frac{1}{g(T)} \cdot \varepsilon x_1$ . Since  $\mathbb{E}[\text{Envy}(A^T)] \leq f(T)$ , we have that  $x_1 \leq \frac{g(T)f(T)}{\varepsilon}$ . In other words, with probability  $1 - \frac{1}{g(T)}$ , every agent receives a number of items within  $[\frac{T}{n^2} - x_1, \frac{T}{n^2} + x_1]$ . Because the first  $\frac{T}{n}$  items are identical for all instances, the bounds on the number of items received from the segment also hold for instances  $I_2, \dots, I_n$ .

Similarly, we consider the behavior of the algorithm faced with  $I_2$  and inspect the end of segment 2. Suppose that, conditioned on the algorithm having allocated each agent a number of items in  $[\frac{T}{n^2} - x_1, \frac{T}{n^2} + x_1]$  from the first segment, some agent  $k$  receives at least  $\frac{T}{n^2} + x_2$  items from the second segment with probability at least  $\frac{1}{g(T)-1}$ . This translates to an (unconditional) probability of at least  $(1 - \frac{1}{g(T)}) \cdot \frac{1}{g(T)-1} = \frac{1}{g(T)}$  and, for similar reasons as before, it follows that

$$x_2 \leq \frac{g(T)f(T)}{\varepsilon} + \frac{2x_1}{\varepsilon} \leq \frac{g(T)f(T)}{\varepsilon} \left(1 + \frac{2}{\varepsilon}\right).$$

Together, after segment 2 we know that with probability  $1 - \frac{2}{g(T)}$ , every agent receives a number of items within  $[\frac{T}{n^2} - x_i, \frac{T}{n^2} + x_i]$  in each segment  $i \in \{1, 2\}$ . We continue this analysis for larger  $i$  with  $x_i = \frac{g(T)f(T)}{\varepsilon} \left(1 + \frac{2}{\varepsilon}\right)^{i-1}$ .

Finally, we analyze efficiency of the algorithm on instance  $I_n$ . Each agent receives expected utility at most  $(1 + (n-1)\varepsilon)(\frac{T}{n^2} + x_n) + \frac{n}{g(T)}T$ , where the additional  $\frac{n}{g(T)}T$  term accounts for the worst case allocation assuming large deviations. Now

$$\frac{u_i}{1/n + \varepsilon} < \frac{T}{n} \cdot \frac{1 + (n-1)\varepsilon}{1 + \varepsilon n} \cdot \left(1 + \frac{n^2}{\varepsilon} \frac{f(T)g(T)}{T} \left(1 + \frac{2}{\varepsilon}\right)^{n-1} + n^3 g(T)\right).$$

For large enough  $T$  satisfying  $\frac{g(T)f(T)}{T} < \frac{1}{2} \frac{\epsilon}{1+(n-1)\epsilon} \cdot \frac{\epsilon}{n^2(1+2/\epsilon)^{n-1}}$  and  $g(T) < \frac{\epsilon}{2n^3}$ , the expected utility for every agent  $i$  is upper bounded by  $\frac{T}{n} \cdot (\frac{1}{n} + \epsilon)$ . On the other hand, the single allocation that gives items  $[\frac{T}{n}(i-1) + 1, \dots, \frac{T}{n}i]$  to agent  $i$  yields utility  $u_i = \frac{T}{n}$  for  $i \in \mathcal{N}$ . We conclude that an allocation algorithm with vanishing envy is not  $(\frac{1}{n} + \epsilon)$ -Pareto efficient for  $\epsilon > 0$ .  $\square$

## 4. Simultaneous Fairness and Efficiency Against Bayesian Adversaries

Having established that it is impossible to simultaneously provide non-trivial fairness and efficiency guarantees against worst-case adversaries, we turn our attention to weaker, Bayesian adversaries.

We start in Section 4.1 with identical agents and independent and identically distributed (i.i.d.) items. Using a result by Dickerson et al. (2014), we find that it is straightforward to simultaneously achieve Pareto efficiency and either envy-freeness with high probability or envy-freeness up to one good.

We then proceed to our main positive result, an algorithm for correlated agents with i.i.d. items that gives the optimal fairness and efficiency guarantees. This, of course, implies the same result for independent agents with i.i.d. items. In Section 4.2, we highlight key insights while ignoring some of the technical obstacles and find an ex post Pareto efficient algorithm achieving the weaker fairness guarantee of vanishing envy. We develop the algorithm fully in Section 4.3.

### 4.1. Identical Agents with IID Items

Suppose an adversary picks a single distribution  $D$ , with support  $G_D$  of size  $m$ , and each  $v_{it}$  is sampled i.i.d. from  $D$ , for all agents  $i \in \mathcal{N}$  and all items  $t \in [T]$ . Consider the following variant of the algorithm discussed in Example 1.

ALGORITHM 1. If  $D$  is a point mass, allocate arriving items in a round-robin manner. Otherwise allocate each item  $t$  to the agent  $i$  with the maximum value  $v_{it}$ , breaking ties uniformly at random.

Efficiency and fairness can be simultaneously achieved using Algorithm 1.

THEOREM 4. *Algorithm 1 outputs an allocation that is always Pareto efficient. Furthermore, for all  $\epsilon > 0$ , there exists  $T_0 = T_0(\epsilon)$ , such that if  $T \geq T_0$ , either the output allocation satisfies pairwise EF1 or it is envy-free with probability at least  $1 - \epsilon$ .*

This result was essentially proved in a different context by [Dickerson et al. \(2014\)](#). They consider a *static* setting with  $T$  items and  $n$  agents where  $v_{it}$  is drawn from a distribution  $D_i$ . It is found that, under mild conditions on the distributions, an envy free allocation exists with probability 1 as  $T \rightarrow \infty$  as long as each agent receives roughly  $T/n$  goods, and each agent has higher expected utility for the goods they are allocated than the rest. We remove these conditions with a slight and unavoidable complication in the fairness guarantee. Full details appear in [Section EC.5.1](#).

## 4.2. Vanishing Envy and Pareto Efficiency for Correlated Agents

Ideally, we would retain the simplicity of [Algorithm 1](#) and extend it to work with stronger adversaries. However, when agents' valuations are no longer identical but merely independent, asking that agent  $i$  has the highest value for an arbitrary item with probability  $1/n$  is a fairly strong requirement, so the result of [Dickerson et al. \(2014\)](#) no longer holds. One possible approach is to assign item  $t$  to the agent  $i$  for whom  $F_{D_i}(v_{it})$  is highest, where  $F_{D_i}$  is the quantile function for agent  $i$ 's value distribution. In fact, this approach is fruitful if one focuses solely on fairness, as shown by [Kurokawa et al. \(2016\)](#). Unfortunately, the resulting allocation is not guaranteed to be Pareto efficient, as the following example shows.

**EXAMPLE 3.** Consider an instance with  $n = 2$  where  $v_{1t} \sim \mathcal{U}[0, 1]$  and  $v_{2t} \sim \mathcal{U}[\frac{1}{2} - \varepsilon, \frac{1}{2} + \varepsilon]$  for all  $t \in \mathcal{G}$ , where  $\mathcal{U}$  denotes the uniform distribution. Agent 2 cares chiefly about how many items they receive. Suppose each item  $t$  is allocated to the agent  $i$  for whom  $F_{D_i}(v_{it})$  is greatest. Roughly, we can construct a Pareto improvement by transferring one item  $t$  for which  $F_{D_2}(v_{2t}) > F_{D_1}(v_{1t}) = 1 - \varepsilon$  from agent 2 to agent 1, and transferring back multiple items for which  $F_{D_2}(v_{2t}) < F_{D_1}(v_{1t}) = \varepsilon$

All in all, achieving fairness and efficiency simultaneously beyond identical agents seems a lot more intricate than either property in isolation. We will skip the independent agent case altogether, and directly study the harder problem of correlated agents: each item  $t$  draws its type  $\gamma_j$  from a distribution  $D$ . Items are i.i.d. but agent values can be correlated.

Before we present the optimal algorithm we illustrate some key ideas by giving a simple algorithm that achieves ex post Pareto efficiency and a weaker notion of fairness, namely vanishing envy with

high probability. Recall that  $f_D(\gamma_j)$  is the probability that an item drawn from  $D$  has type  $\gamma_j$ ,  $G_D$  is the support of  $D$ ,  $|G_D| = m$  and  $v_i(\gamma_j)$  is the value of an item of type  $\gamma_j$  to agent  $i$ . For ease of notation we sometimes refer to item type  $\gamma_j$  as  $j$ .

Our approach is to solve an offline *divisible* item allocation problem as an intermediate step. The resulting fractional allocation is  $X \in [0, 1]^{n \times m}$ , where  $n$  is the number of agents and  $m = |G_D|$  the number of types of items in the support of  $D$ . For each  $i \in \mathcal{N}$ ,  $X_{ij} \in [0, 1]$  is the proportion of item type  $j$  allocated to agent  $i$ .  $X$  is constrained to be feasible, i.e.,  $\sum_{i \in \mathcal{N}} X_{ij} = 1$  for all types  $j \in G_D$ . The  $i^{\text{th}}$  row of  $X$ , denoted  $X_i$ , is the fractional allocation received by agent  $i \in \mathcal{N}$ .

ALGORITHM 2: PARETO EFFICIENT ROUNDING

Input: Distribution  $D$  over item types, agent valuation functions  $v_i$ .

1. For each  $\gamma_j \in G_D$  and  $i \in \mathcal{N}$ , set  $v'_i(\gamma_j) = v_i(\gamma_j)f_D(\gamma_j)$ .
2. Find the divisible allocation  $X$  of  $G_D$  that maximizes the product of utilities with respect to  $v'$ .
3. In the online setting, allocate the newly arrived item  $t$  with type  $\gamma_j$  to agent  $i$  with probability  $X_{ij}$ , for all  $t = 1, \dots, T$ .

We first show that Algorithm 2 always produces a Pareto efficient allocation. In fact, we show something much stronger: *every* rounding of *every* Pareto efficient fractional allocation  $X$  results in an ex post Pareto efficient allocation.

**THEOREM 5.** *Given a distribution  $D$  over  $m$  item types and valuation function  $v_i$  for each agent  $i \in \mathcal{N}$ , let  $X$  be a Pareto efficient allocation of  $G_D$  under valuation functions  $v'_i$ , with  $v'_i(\gamma_j) = v_i(\gamma_j) \cdot f_D(\gamma_j)$ . Let  $S$  be a set of  $T$  items drawn from  $D$ , and  $A = (A_1, \dots, A_n)$  any allocation of  $S$  where an item of type  $\gamma_j$  is allocated to agent  $i$  only if  $X_{ij} > 0$ . Then  $A$  is Pareto efficient under  $v$ .*

*Proof.* By definition  $v'_i$  is  $v_i$  scaled by the probability  $f_D(\gamma_j)$  that type  $\gamma_j$  appears. Let  $\tilde{v}_i$  be the valuation function when scaling with respect to the observed frequencies in  $S$ , i.e.  $\tilde{v}_i(\gamma_j) = v_i(\gamma_j) \cdot fr(\gamma_j)$ , where  $fr(\gamma_j) = \sum_{t \in S} \mathbf{1}\{\text{item } t \text{ has type } \gamma_j\}$ . We prove the theorem in two steps. First, we show that  $X$  is Pareto efficient under  $\tilde{v}$ . Second, we show this implies  $A$  is Pareto efficient under  $v$ .

Suppose for contradiction that  $X$  is not Pareto efficient under  $\tilde{v}$ . Then there exists an allocation  $X'$  that dominates  $X$  under  $\tilde{v}$ . Let  $\Delta = X' - X$  denote the of item transfers needed to go from  $X$  to  $X'$ . For all  $c \in [0, 1]$ , the allocation  $X + c\Delta$  is feasible and still dominates  $X$  under  $\tilde{v}$ . We construct  $\Delta'$ , where  $\Delta'_{ij} = \Delta_{ij} \cdot fr(\gamma_j)/f_D(\gamma_j)$ . Observe that the change in utilities induced by transfers  $\Delta'$  under  $v'$  equals the change in utilities induced by transfers  $\Delta$  under  $\tilde{v}$ . Therefore, the (possibly infeasible) allocation  $X + \Delta'$  dominates  $X$  under  $v'$ , as does  $X + c\Delta'$  for all  $c \in [0, 1]$ .

Consider  $X + c\Delta'$  for  $0 < c = \min_k f_D(\gamma_k)/fr(\gamma_k)$ . Notice  $c$  is well-defined and  $(X + c\Delta')_{ij} = \delta_j X'_{ij} + (1 - \delta_j)X_{ij} \in [0, 1]$ , where  $\delta_j = c \cdot fr(\gamma_j)/f_D(\gamma_j) \leq 1$ . We conclude  $X + c\Delta'$  is feasible and dominates  $X$  under  $v'$ , a contradiction.

Next, we show that if  $X$  is Pareto efficient under  $\tilde{v}$ , then  $A$  is Pareto efficient under  $v$ . Suppose that  $A$  is not efficient under  $v$  and is dominated by an allocation  $A'$ . Let  $Y, Y'$  be fractional allocations of  $G_D$ , where  $Y_{ij} = (\sum_{t \in [T]} \mathbb{1}\{t \in A_i \text{ and item } t \text{ has type } \gamma_j\}) / (\sum_{t \in [T]} \mathbb{1}\{\text{item } t \text{ has type } \gamma_j\})$  is the fraction of items of type  $\gamma_j$  given to agent  $i$  in  $A$ . Define  $Y'$  similarly for  $A'$ .

The utility of agent  $i$  receiving allocation  $Y$  under  $\tilde{v}$  is:

$$\begin{aligned} \sum_{j \in G_D} \tilde{v}_i(\gamma_j) Y_{ij} &= \sum_{j \in G_D} v_i(\gamma_j) \cdot fr(\gamma_j) \cdot \frac{\sum_{t \in [T]} \mathbb{1}\{t \in A_i \text{ and item } t \text{ has type } \gamma_j\}}{\sum_{t \in [T]} \mathbb{1}\{\text{item } t \text{ has type } \gamma_j\}} \\ &= \sum_{j \in G_D} v_i(\gamma_j) \cdot \left( \sum_{t \in [T]} \mathbb{1}\{\text{item } t \text{ has type } \gamma_j\} \right) \cdot \frac{\sum_{t \in [T]} \mathbb{1}\{t \in A_i \text{ and item } t \text{ has type } \gamma_j\}}{\sum_{t \in [T]} \mathbb{1}\{\text{item } t \text{ has type } \gamma_j\}} \\ &= \sum_{t \in [T]} v_{it} \cdot \mathbb{1}\{g_t \in A_i\}, \end{aligned}$$

i.e., the same as for allocation  $A$  under  $v$ . Similarly with  $A', Y'$ . Let  $\Delta = Y' - Y$ . For any  $c > 0$ ,  $c\Delta$  is a Pareto improvement on any allocation under  $\tilde{v}$ , and therefore, the (potentially infeasible) allocation  $X + c\Delta$  dominates  $X$  under  $\tilde{v}$ . In EC.5.2 we show how to find  $c^* > 0$  such that  $X + c^*\Delta$  is feasible. Combining the two steps completes the proof.  $\square$

Maximizing the product of utilities leads to a fractional Pareto efficient allocation. Therefore, Theorem 5 implies that Algorithm 2 is ex post Pareto efficient. We now show it also guarantees a notion of fairness slightly weaker than vanishing envy, namely vanishing envy with high probability.

**THEOREM 6.** *For all  $\varepsilon > 0$ , there exists  $T_0 = \sqrt{4/\varepsilon}$ , such that if  $T \geq T_0$ , Algorithm 2 outputs an allocation  $A$  such that for all agents  $i, j$ ,  $\text{Envy}_{i,j}(A) \in o(T)$  with probability at least  $1 - \varepsilon$  and  $\mathbb{E}[\text{Envy}_T] \in O(\sqrt{T \log T})$ .*

Proof of Theorem 6. The fractional allocation  $X$  that maximizes the product of utilities is envy-free (Varian 1974), which implies  $\sum_{k \in [m]} v_i(\gamma_k) f_D(\gamma_k) X_{ik} \geq \sum_{k \in [m]} v_i(\gamma_k) f_D(\gamma_k) X_{jk}$  for all pairs of agents  $i, j \in \mathcal{N}$ .

Let  $A$  be the allocation that results from Algorithm 2. Agent  $i$ 's value for agent  $j$ 's bundle,  $v_i(A_j)$ , is a random variable that depends on randomness in both the algorithm and item draws. Let  $I_t^{k,j}$  be an indicator random variable for the event that item  $t$  is of type  $\gamma_k$  and is assigned to agent  $j$ . For any pair of agents  $i, j \in \mathcal{N}$ ,  $v_i(A_j) = \sum_{t \in [T]} \sum_{k \in [m]} v_i(\gamma_k) I_t^{k,j}$ . Therefore,  $\mathbb{E}[v_i(A_j)] = T \cdot \sum_{k \in [m]} v_i(\gamma_k) f_D(\gamma_k) X_{jk}$ . By the envy-freeness of the fractional allocation,  $\mathbb{E}[v_i(A_i)] \geq \mathbb{E}[v_i(A_j)]$ .

It now follows from Hoeffding's inequality (Hoeffding 1963) with parameter  $\delta = \sqrt{T \log T}$  that

$$\Pr \left[ v_i(A_i) - \mathbb{E}[v_i(A_i)] \leq -\sqrt{T \log T} \right] \leq 2 \exp \left( -\frac{2T \log T}{T} \right) = \frac{2}{T^2}.$$

Similarly we bound the deviation of  $v_i(A_j)$ ,  $\Pr[v_i(A_j) - \mathbb{E}[v_i(A_j)] \geq \sqrt{T \log T}] \leq 2/T^2$ . Together, we conclude for  $T_0 = \sqrt{4/\varepsilon}$  that  $\text{Envy}_{i,j}(A) = \max\{v_i(A_j) - v_i(A_i), 0\} \leq 2\sqrt{T \log T} \in o(T)$  with probability at least  $1 - \frac{4}{T^2} \geq 1 - \varepsilon$ .

To compute expected envy at  $T$ , we set  $\varepsilon = \frac{1}{n^2 T}$  in the above, observe  $T_0 = 2n\sqrt{T} < T$  as required, condition on some pairwise envy exceeding  $2\sqrt{T \log T}$  and apply a union bound to obtain

$$\mathbb{E}[\text{Envy}_T] \leq \sum_{i,j \in \mathcal{N}} \Pr[\text{Envy}_{i,j} \geq 2\sqrt{T \log T}] \cdot T + 1 \cdot 2\sqrt{T \log T} = n^2 T \cdot \frac{1}{n^2 T} + 2\sqrt{T \log T} \in O(\sqrt{T \log T}) \quad \square$$

### 4.3. Beyond Vanishing Envy: Optimal Fairness for Correlated Agents

In the proof of Theorem 6 we use standard tail inequalities to show that, with high probability, the envy between any two agents does not deviate from its expectation by more than  $O(\sqrt{T \log T})$ . The divisible allocation is envy-free, and rounding it online leads to vanishing envy. If, instead, the divisible allocation  $X$  used to guide the online decisions satisfied *strong envy freeness*, so for every pair of agents  $i, j \in \mathcal{N}$ ,  $v_i(X_i) > v_i(X_j)$ , then we could argue similarly that the online integral allocation would be envy-free with high probability.

Unfortunately, strong envy-free allocations do not always exist, even for divisible items, as in the case of two agents with identical valuation functions. Interestingly, this condition is also sufficient:



as long as no two agents have identical valuation functions (up to multiplicative factors), a strongly envy-free allocation exists (Barbanel 2005). However, this is no longer sufficient if we want both Pareto efficiency and strong envy-freeness (see EC.6.1 for an example).

Nevertheless, we can achieve a notion of fairness offline that is weaker than strong envy-freeness, but sufficient for our purposes. We say agent  $i$  is *indifferent* to agent  $j$  if  $v_i(X_i) = v_i(X_j)$ . We construct a directed *indifference graph*  $I(X)$  with a vertex for each agent  $i \in \mathcal{N}$  and containing edge  $(i, j)$  exactly when  $i$  is indifferent to  $j$  under  $X$ . For an envy-free allocation  $X$ , the absence of  $(i, j)$  in  $I(X)$  implies that  $v_i(X_i) > v_i(X_j)$ , i.e., strong (pairwise) envy-freeness. We consider the following notion of fairness.

**DEFINITION 1.** A fractional allocation  $X$  is *clique identical strongly envy-free* (CISEF) if (1)  $X$  is envy-free; (2)  $I(X)$  is a disjoint union of cliques; (3) agents in the same clique have identical valuations (up to a multiplicative factor) for all items allocated to any member of the clique; and (4) agents in the same clique have identical allocations.

Our main structural result is that, though Pareto efficiency is incompatible with strong envy-freeness, it is compatible with CISEF.

**THEOREM 7.** *Given any instance with  $m$  divisible items and  $n$  additive agents, there always exists an allocation that is simultaneously clique identical strongly envy free (CISEF) and Pareto efficient.*

This result is constructive and somewhat technical (see Section EC.6.2). We provide a sketch.

**Proof sketch for Theorem 7.** We start by solving the Eisenberg-Gale convex program (henceforth the E-G program) or, equivalently, by finding the competitive equilibrium from equal incomes. This is a standard approach for finding an envy-free and Pareto efficient allocation. Recall that the E-G program with “budgets”  $\mathbf{e}$  consists of

$$\max_X \sum_{i=1}^n e_i \log \sum_{j=1}^m v_{ij} X_{ij}, \quad \text{subject to} \quad \sum_{i=1}^n X_{ij} \leq 1, \forall j \in [m], \quad \text{and} \quad X_{ij} \geq 0, \forall i \in [n], j \in [m].$$

Specifically, we give each agent a budget  $e_i = 1$ , and find market-clearing prices (a price  $p_j$  for each item  $j$ ) such that each agent  $i$  only buys items that maximize her “bang-per-buck” ratio  $v_{ij}/p_j$ . Let  $X_0$  be this initial allocation, and  $\mathbf{p}$  and  $\mathbf{e}$  be the prices and budgets.

Then, at a high level, we proceed by repeatedly altering  $X, \mathbf{p}$  and  $\mathbf{e}$  in such a way that  $X, \mathbf{p}$  remains an optimal solution to the convex program with budgets  $\mathbf{e}$ , while preserving envy-freeness. This terminates when  $X$  satisfies the desired properties. More specifically, at termination  $I(X)$  will be a disjoint set of cliques, where agents in a clique have identical allocations.  $\square$

It is worth highlighting a connection between the indifference graph and the bipartite maximum bang per buck (MBB) graph. Properties of MBB graphs have been crucial to recent algorithmic progress in approximating Nash social welfare (Cole and Gkatzelis (2015), Garg et al. (2018), Chaudhury et al. (2018)) and computing equilibria in Arrow-Debreu exchange markets (Garg and Végh (2019)). In the indifference graph there is an edge from  $i$  to  $j$  when  $i$  is indifferent between her allocation and the allocation of agent  $j$ . We show (Lemma EC.2) this condition is similar to the condition for edges existing in the MBB graph but are unaware of further overlap.

Algorithm 3 is a slightly modified version of Algorithm 2 which, when using a Pareto efficient and CISEF fractional allocation to guide the online allocations, yields an integral allocation that is Pareto efficient ex post and achieves the target fairness properties.

**ALGORITHM 3: PARETO EFFICIENT CLIQUE ROUNDING**

Input: Item distribution  $D$ , agent valuation functions  $v_i$ .

1. For each  $\gamma_j \in G_D$  and  $i \in \mathcal{N}$ , define  $v'_i(\gamma_j) = v_i(\gamma_j)f_D(\gamma_j)$ .
2. Compute a fractional allocation  $X^*$  of  $G_D$  that is Pareto efficient and CISEF under  $v'_i$ . Let  $C_1, \dots, C_s$  be the disjoint cliques of  $I(X^*)$ .
3. In the online setting, assign the newly arrived item  $t$  with type  $\gamma_j$  to clique  $C_i$  with probability  $\sum_{k \in C_i} X_{kj}^*$ . When an item is assigned to a clique  $C_i$ , allocate it to the agent in  $C_i$  who has received the least value so far according to (*all*) agents in the clique.

An algorithm for constructing  $X^*$  can be found in Section EC.6.2. Notice in Algorithm 3 step 3 that there is a unique agent with smallest value (up to tie-breaking) since all agents agree on the value of all items that have gone to the clique (up to multiplicative factors).

**THEOREM 8.** *The allocation resulting from Algorithm 3 is ex post Pareto efficient and, for all  $\varepsilon > 0$ , there exists  $T_0 = T_0(\varepsilon)$ , such that if  $T \geq T_0$ , for all agents  $i, j \in \mathcal{N}$ , either  $i$  envies  $j$  by at most one item with probability 1 or  $i$  does not envy  $j$  with probability at least  $1 - \varepsilon$ .*

Proof of Theorem 8. Let  $X^*$  be the fractional CISEF and Pareto efficient allocation and  $A$  the integral allocation produced by Algorithm 3. Pareto efficiency of  $A$  follows directly from Theorem 5.

For any two agents  $i, j \in \mathcal{N}$ , there are two cases. Suppose  $i$  and  $j$  belong to the same clique  $C_k$ . Let  $S$  be the set of items assigned to  $C_k$  during the execution of Algorithm 3, i.e.  $S = \cup_{\ell \in C_k} A_\ell$ . Agents in  $C_k$  have identical valuations up to a multiplicative factor for the items that they get with positive probability. Therefore, giving each item to the agent that has received the least value so far (according to any agent, as they rank allocations of  $S$  in the same order) ensures that  $\text{Envy}_{i,j}(A) \leq \max_{s \in S} v_{is} \leq 1$ .

Now suppose  $i$  and  $j$  belong to different cliques  $C_i$  and  $C_j$ , respectively. By the definition of a CISEF allocation, we know that  $v_i(X_i^*) = v_i(X_j^*) + c$  for some constant  $c > 0$ .

Let  $\tilde{A}$  be the fractional allocation where every agent  $p$  in clique  $C_p$  receives  $1/|C_p|$  fraction of the items assigned to  $C_p$ . In particular, all  $i' \in C_i$  receive  $\tilde{A}_{i't} = \frac{1}{|C_i|} \mathbb{1}\{t \in A_k \text{ for some agent } k \in C_i\}$ . Similarly for  $\tilde{A}_j$ .  $\tilde{A}_i$  is the average allocation of agents in  $C_i$  (in  $A$ ) and, as argued earlier, the maximum envy for two agents in the same clique is at most 1 in  $A$ . It follows that  $|v_i(A_i) - v_i(\tilde{A}_i)| \leq 1$ . Furthermore, agents in the same clique receive identical allocations in  $\tilde{A}$ , so  $\mathbb{E}[v_i(\tilde{A}_i) - v_i(\tilde{A}_j)] = T\mathbb{E}[v_i(X_i^*) - v_i(X_j^*)] = cT$ .

It follows from Hoeffding's inequality (Hoeffding 1963) that, with probability at least  $1 - \Theta(1/T^2) \geq 1 - \varepsilon$ ,  $v_i(\tilde{A}_j) - v_i(\tilde{A}_i) < 2\sqrt{T \log T} - cT$ . This is negative for sufficiently large  $T$ , so we can pick  $T$  such that  $v_i(\tilde{A}_j) - v_i(\tilde{A}_i) < -2$  with high probability.

We conclude  $v_i(A_j) - v_i(A_i) < v_i(\tilde{A}_j) - v_i(\tilde{A}_i) + 2 < 0$  with probability at least  $1 - \varepsilon$ .  $\square$

## 5. Discussion

We finish with a discussion of several pertinent issues that have not been addressed so far.

*Additivity assumption.* We assume agents have additive valuations. This common assumption is often considered strong. However, for the purpose of defining envy in our online setting we believe it is quite natural. Since items arrive over time and are potentially perishable (as in food bank applications) they are likely used independently of each other. Furthermore, we can reformulate  $\text{Envy}_{ij}(A) = \sum_{t \in A_i} v_{it} - \sum_{t \in A_j} v_{it} = \sum_{t=1}^T v_{it}(\mathbb{I}_{t \in A_i} - \mathbb{I}_{t \in A_j})$  as the sum of per-round envies, so assuming additive valuations amounts to assuming envy is additive over time.

*Computational considerations.* Theorem 7 ensures all our algorithms run in polynomial time. We require an exact solution to the E-G program, which is obtainable in strongly polynomial time (Orlin 2010). The edge-elimination steps happen  $O(n^2)$  times. The only remaining question is the number of bits in the solution  $(X, \mathbf{p})$  and budgets  $\mathbf{e}$ , as the item transfers in Lemmas EC.4 and EC.5 can both increase the length (in bits) of  $X$  and  $\mathbf{e}$ . However, as we discuss in section EC.6.2, this increase is limited to a constant number of bits.

*Future directions.* We very nearly have a complete picture of what is possible when optimizing fairness or efficiency in isolation. The one exception is minimizing fairness for the non-adaptive worst-case adversary, where vanishing envy is certainly possible (the  $\tilde{O}(\sqrt{T/n})$  guarantee of Theorem 1 applies) but we do not even have a super-constant lower bound. An open technical question is what happens when the distributions chosen by the adversary are allowed to depend on  $T$ . Finally, there is a legitimate question of whether it is reasonable to assume perfect information about agent utilities. It may be more realistic to assume partial access to utilities, for example in the form of pairwise comparisons between the item under consideration and previously allocated ones.

## Details omitted from the main text

### EC.1. Asymptotic Notation

Let  $f(n), g(n)$  be functions defined on the natural numbers. We say  $f(n) \in O(g(n))$  when there exists constants  $c$  and  $n_0$  so that  $f(n) \leq c \cdot g(n)$  for all  $n \geq n_0$ . Informally we can say  $f$  grows no faster than  $g$ . We say  $f(n) \in \Omega(g(n))$  when  $g(n) \in O(f(n))$ , and  $f(n) \in \Theta(g(n))$  when  $g(n) \in O(f(n))$  and  $f(n) \in O(g(n))$ . The strict version of this relationship is indicated as  $f(n) \in o(g(n))$ , which means that  $f(n) \in O(g(n))$  but  $f(n) \notin \Theta(g(n))$ . Finally  $f(n) \in \omega(g(n))$  when  $g(n) \in o(f(n))$ .

### EC.2. Improving on “EF1 or EF w.h.p” is Impossible

In Section 4 we show that it is possible to achieve Pareto efficiency and be either envy-free up to 1 good, or envy-free with high probability. We now show that it is impossible to improve this fairness guarantee against the weakest (Bayesian) adversary. Recall that this adversary specifies a distribution  $D^V$  and item values  $v_{it}$  are drawn i.i.d. from  $D^V$ .

First, we show that it is not always possible to guarantee envy-freeness with high probability. Define  $D^V$  to be the uniform distribution over the set  $\{1\}$ . Note that whenever  $T$  is not a multiple of  $n$ , the allocation will not be envy-free. Next, we show that for any  $x$ , envy-freeness up to  $x$  goods, is not an achievable guarantee. We use the construction of the lower bound in Section 3.2, which assumes item values bounded within  $[0, 1]$ . Recall the following result:

**THEOREM 2.** *For  $n \geq 2$  and  $r < 1$ , there exists an adversary strategy for setting item values such that any algorithm must have  $\text{Envy}(A) \in \Omega((T/n)^{r/2})$ , where  $A$  is the allocation  $T$  items.*

In the proof of Theorem 2, an adversary strategy is specified where all agents, excluding the first two, have no value for any item. The value of the arriving item to the first two agents depends solely on a state machine, which is described fully in Section 3.2.

Set  $r = \frac{1}{2}$ . Let  $c, T_0$  to be the constants such that for any  $T \geq T_0$ , the adversary can guarantee  $\text{Envy}(A) \geq c(T/n)^{r/2} = c(T/n)^{1/4}$ . We then take any  $T' \geq T_0$  where  $c(T'/n)^{1/4} > x$  and set  $D$  to be the uniform distribution over  $\{0, 1, \nu_1, \dots, \nu_{T'}\}$ .

We claim that for any  $T \geq T'$ , there is a positive probability that the allocation will not be envy-free up to  $x$  goods. First, for the given algorithm, there is a positive probability that the first  $T$  items drawn will follow the adversary strategy. Then, for the remaining items, there is a positive probability that all the items will have no value to any agent. Thus, for the final allocation  $A$ , we will have  $\text{Envy}(A) > x$ . Finally, observe that when item values are bounded within  $[0, 1]$ ,  $\text{Envy}(A) > x$  implies that  $A$  is not envy-free up to  $x$  goods.

### EC.3. Details and Proofs from Section 3.1.

We first take a closer look at the game tree with nodes on  $T + 1$  levels. Every node on level  $1, \dots, T$  has  $n$  outgoing arcs labeled  $1, \dots, n$ . The leaf nodes on level  $T + 1$  are labeled by the maximum envy for the corresponding path, which defines an allocation of the  $T$  items. Let  $\Omega$  be the set of all paths from the root to a leaf node, so  $|\Omega| = n^T$ . Equivalently,  $\Omega$  is the set of all possible allocations of the  $T$  items. For an allocation  $\omega \in \Omega$ , denote by  $\omega_t \in [n]$  the agent to whom item  $t \in [T]$  was allocated by  $\omega$ .

A fully adaptive strategy  $s$  for the adversary is defined by labeling every internal node  $u$  with a value vector  $s(u)$ , where  $s(u)_i$  is the value of agent  $i$  for the item corresponding to node  $u$ . The algorithm's strategy consists of selecting an outgoing edge, corresponding to an allocation of the item with valuation  $s(u)$ , at every node  $u$ . The adversary's strategy is allowed to depend on the allocations and valuations so far, i.e., the path from the root to  $u$ .

For a given adversary strategy  $s$  and an allocation  $\omega$ , let  $\text{Envy}^{ij}(s, \omega)$  denote the envy of agent  $i$  for agent  $j$ . Denote with  $\text{Envy}(s, \omega) = \max_{i, j \in [n]} \text{Envy}^{ij}(s, \omega)$  the maximum envy experienced by any agent. The objective of the adversary is to choose a strategy  $s$  that maximizes the expected envy  $\mathbb{E}[\text{Envy}(s, \omega)]$ , where the expectation is taken over allocating every item uniformly at random.

Consider the algorithm that allocates every item uniformly at random. This is equivalent to picking a random outgoing edge at each node  $u$ .

**LEMMA 1** *The adversary has an optimal adaptive strategy that labels every internal node of the game tree with a vector in  $\{0, 1\}^n$ .*

Proof of Lemma 1. Assume for the sake of contradiction that the adversary does not have an optimal strategy which assigns integral vectors to the nodes of the (adversary-centric) game tree. Let  $s$  be the optimal strategy with the smallest number of fractional values. Without loss of generality, let  $u$  be a node on layer  $\ell \in [T]$  for which the value assigned to player  $i \in [n]$  is fractional, i.e.,  $0 < s(u)_i < 1$ . The values  $\ell$  and  $i$  are fixed for the remainder of this proof. Define alternative strategies  $s'$  and  $s''$  identical to  $s$ , except that  $s'(u)_i = 1$  and  $s''(u)_i = 0$ . We wish to arrive at the contradiction that  $\mathbb{E}[\text{Envy}(s, \omega)] \leq \mathbb{E}[\text{Envy}(s^*, \omega)]$  for  $s^* = s'$  or  $s''$ , where the expectation is over the randomness of the allocation algorithm. Denote with  $\Omega_u$  all paths passing through  $u$ . The envy associated with paths in  $\Omega \setminus \Omega_u$  is unaffected by the move from  $s$  to  $s'$  or  $s''$  and may be safely ignored.

When agent  $i$  is not the unique agent with maximum envy, it holds that  $\text{Envy}(s, \omega) \leq \text{Envy}(s', \omega)$  and  $\text{Envy}(s, \omega) \leq \text{Envy}(s'', \omega)$  as desired (recall that changing agent  $i$ 's valuation for an item does not affect other agents' envy). It remains to consider the set of paths

$$\Omega_u^+ = \left\{ \omega \in \Omega : \max_{j \in [n]} \text{Envy}^{ij}(s, \omega) > \max_{j \in [n] \setminus \{i\}} \max_{k \in [n]} \text{Envy}^{jk}(s, \omega) \right\},$$

in which agent  $i$  is the unique agent with maximum envy (and this envy is strictly positive). We can further partition  $\Omega_u^+$  according to which agent receives item  $\ell$ ; let  $\Omega_u^{+,j}$  be the set of paths in  $\Omega_u^+$  in which agent  $j \in [n]$  gets item  $\ell$ , and for any  $J \subseteq [n]$ , set  $\Omega_u^{+,J} = \cup_{j \in J} \Omega_u^{+,j}$ . We analyze three different cases: (1) whether the player that gets item  $\ell$  is player  $i$ , (2) a player  $j^*$  for whom player  $i$  has maximum envy, or (3) another player. Define

$$J^* = \left\{ j^* \in [n] : \text{Envy}^{ij^*}(s, \omega) = \max_{j \in [n]} \text{Envy}^{ij}(s, \omega) \right\}.$$

Also, for convenience, set  $f = s(u)_i$  and  $J^< = [n] \setminus \{J^* \cup \{i\}\}$ .

We first look at  $s'$ . The three cases are:

1. For  $\omega \in \Omega_u^{+,i}$ :  $\text{Envy}(s, \omega) - (1 - f) \leq \text{Envy}(s', \omega) \leq \text{Envy}(s, \omega)$ .
2. For  $\omega \in \Omega_u^{+,J^*}$ :  $\text{Envy}(s', \omega) = \text{Envy}(s, \omega) + (1 - f)$ .
3. For  $\omega \in \Omega_u^{+,J^<}$ :  $\text{Envy}(s, \omega) \leq \text{Envy}(s', \omega) \leq \text{Envy}(s, \omega) + (1 - f)$ .

The only outcomes where envy can decrease when changing the adversary's strategy from  $s$  to  $s'$  are those in  $\Omega_u^{+,i}$ . We can compute the effect of changing  $s$  to  $s'$  on the expected maximum envy as

$$\begin{aligned}
\mathbb{E}[\text{Envy}(s, \omega)] &= \sum_{\omega \in \Omega} \Pr[\omega] \cdot \text{Envy}(s, \omega) \\
&= \frac{1}{n^T} \left( \sum_{\omega \in \Omega_u^{+,i}} \text{Envy}(s, \omega) + \sum_{\omega \in \Omega_u^{+,J^*}} \text{Envy}(s, \omega) + \sum_{\omega \in \Omega_u^{+,J^<}} \text{Envy}(s, \omega) \right) \\
&\leq \frac{1}{n^T} \left( \sum_{\omega \in \Omega_u^{+,i}} (\text{Envy}(s', \omega) + (1-f)) + \sum_{\omega \in \Omega_u^{+,J^*}} (\text{Envy}(s', \omega) - (1-f)) + \sum_{\omega \in \Omega_u^{+,J^<}} \text{Envy}(s', \omega) \right) \\
&= \mathbb{E}[\text{Envy}(s', \omega)] + \frac{1-f}{n^T} (|\Omega_u^{+,i}| - |\Omega_u^{+,J^*}|).
\end{aligned}$$

If  $|\Omega_u^{+,i}| \leq |\Omega_u^{+,J^*}|$ , it follows that  $\mathbb{E}[\text{Envy}(s, \omega)] \leq \mathbb{E}[\text{Envy}(s', \omega)]$ . Assume therefore that  $|\Omega_u^{+,i}| > |\Omega_u^{+,J^*}|$ . An identical analysis for  $s''$  shows that

1. For  $\omega \in \Omega_u^{+,i}$ :  $\text{Envy}(s'', \omega) = \text{Envy}(s, \omega) + f$ .
2. For  $\omega \in \Omega_u^{+,J^*}$ :  $\text{Envy}(s, \omega) - f \leq \text{Envy}(s'', \omega) \leq \text{Envy}(s, \omega)$ .
3. For  $\omega \in \Omega_u^{+,J^<}$ :  $\text{Envy}(s, \omega) = \text{Envy}(s'', \omega)$ .

Expanding the computation of the expected value as before shows

$$\mathbb{E}[\text{Envy}(s, \omega)] \leq \mathbb{E}[\text{Envy}(s', \omega)] + \frac{f}{n^T} (-|\Omega_u^{+,i}| + |\Omega_u^{+,J^*}|).$$

By assumption  $|\Omega_u^{+,i}| > |\Omega_u^{+,J^*}|$ , so  $\mathbb{E}[\text{Envy}(s, \omega)] \leq \mathbb{E}[\text{Envy}(s'', \omega)]$ , concluding the proof.  $\square$

**LEMMA 2** *The adversary has an optimal adaptive strategy that labels every internal node of the game tree with the vector  $\mathbf{1}^n$ .*

*Proof of Lemma 2.* By Lemma 1, the adversary has an optimal strategy that labels every internal node with a vector in  $\{0, 1\}^n$ . Let  $s$  be such an optimal strategy with the smallest number of zeros, and suppose (for the sake of contradiction) that there exist internal nodes that are not labeled  $\mathbf{1}^n$ . Let  $u$  on layer  $\ell \in [T]$  be the node closest to a leaf node for which  $s(u)$  contains a 0 and  $s(u') = \mathbf{1}^n$  for all descendants  $u'$  of  $u$ . Without loss of generality assume  $s(u)_i = 0$ , so agent



$i$  has value 0 for item  $\ell$  at node  $u$ . Define a strategy  $s'$  identical to  $s$  except that  $s'(u)_i = 1$ . Let  $j(\omega) \in \arg \max_{j \in [n]} \text{Envy}^{ij}(s, \omega)$ .

For any fixed  $\omega \in \Omega$ , changing  $s$  to  $s'$  only changes the envy of agent  $i$  and only for paths that go through  $u$ . In particular, if  $\omega_\ell \neq i$ , the envy of agent  $i$  toward agent  $\omega_\ell$  increases by 1, which only helps the adversary. By contrast, if  $\omega_\ell = i$ , the envy of agent  $i$  decreases by 1, toward every agent  $j$  such that  $\text{Envy}^{ij}(s, \omega) > 0$ ; the maximum envy,  $\text{Envy}(s, \omega)$ , is only affected if  $\text{Envy}(s, \omega) = \text{Envy}^{i, j(\omega)}(s, \omega)$ .

Thus, let  $\omega \in \Omega$  be an arbitrary path going through  $u$  with  $\omega_\ell = i$  and satisfying  $\text{Envy}(s, \omega) = \text{Envy}^{i, j(\omega)}(s, \omega) > 0$ . Since agent  $i$  may not have been the unique agent having envy equal to  $\text{Envy}(s, \omega)$ ,  $\text{Envy}(s', \omega) \geq \text{Envy}(s, \omega) - 1$ . Now consider the path  $\omega'$  that is identical to  $\omega$  except that  $\omega_\ell = j(\omega)$ . Observe that  $\text{Envy}(s', \omega') = \text{Envy}(s, \omega') + 1$ . Hence, any decrease in envy due to allocating item  $\ell$  to agent  $i$  on  $\omega$  is compensated for (in the calculation of expected envy) along  $\omega'$ . Since  $\omega$  was picked arbitrarily and the mapping  $\omega \mapsto \omega'$  is injective, it follows that the expected envy under  $s'$  is at least the expected envy under  $s$ , and  $s'$  has fewer zeros than  $s$ , contradicting our assumption on  $s$ .  $\square$

The version of Bernstein's inequality used is:

LEMMA EC.1 (**Bernstein 1946**). *Let  $X_1, \dots, X_T$  be independent variables with  $\mathbb{E}[X_t] = 0$  and  $|X_t| \leq M$  almost surely for all  $t \in [T]$ . Then, for all  $\lambda > 0$ ,*

$$\Pr \left[ \sum_{t=1}^T X_t > \lambda \right] \leq \exp \left( - \frac{\frac{1}{2} \lambda^2}{\sum_{t=1}^T \mathbb{E}[X_t^2] + \frac{1}{3} M \lambda} \right).$$

THEOREM 2. *For any  $n \geq 2$  and  $r < 1$ , there exists an adversary strategy for setting item values such that any algorithm must have  $\text{Envy}_T \in \Omega((T/n)^{r/2})$ .*

Proof of Theorem 2. We augment the instance of Figure 1 in the following way. In addition to the first two agents,  $L$  and  $R$ , we have  $n - 2$  other agents. Each of these other agents will not value *any* of the items that arrive; hence, the nonzero values remain the same as before. State transitions work as follows. If the algorithm allocates an item to agent  $L$  or agent  $R$ , the transitions are the same as when  $n = 2$ . Otherwise, the adversary will remain in the same state.

Let  $T_0$  be the number of items allocated to either agent  $L$  or  $R$ . We break the analysis into two cases. First, if  $T_0 \in \Omega(T/n)$ , then,  $\text{Envy}_T \in \Omega((T/n)^{r/2})$  by the analysis of Lemma 3. Otherwise,  $T_0 \in o(T/n)$  and therefore  $T - T_0 \in \Theta(T)$ , i.e., agents 3 through  $n$  receive many items. This implies that there exists an agent  $i \in [3, n]$  that is allocated  $\Omega(T/n)$  items. Without loss of generality, at least half these items were allocated in the left states, in which agent  $L$  values each item at 1, so that agent  $L$  has  $\Omega(T/n)$  value for the items received by agent  $i$ . The value of agent  $L$  for her own allocation is at most  $O(T_0)$ , i.e.,  $o(T/n)$ . Therefore, the envy of agent  $L$  for agent  $i$  is at least  $\Theta(T/n) - o(T/n) \in \Theta(T/n)$ .  $\square$

#### EC.4. Proofs from Section 3.2.

LEMMA 4.  $\mathcal{A}^*(K) \in \text{OPT}(K)$ .

Proof of Lemma 4. An algorithm that starts at state 0 and spends  $T'$  steps in the left states can be described as a sequence of choices  $s_t \in \{L, R\}$  for  $t \in [T']$  such that  $s_1 = L$ , and at every  $t \in [T']$ , agent  $L$  has received at least as many of the first  $t$  items as agent  $R$  (to avoid entering the right states). We refer to the state at time  $t$  as the state *after* the algorithm choice  $s_t$ .

Consider any  $\mathcal{A}(K) \in \text{OPT}(K)$ . We show that the corresponding sequence of allocations satisfy: (1) at time  $T'$  the state is  $L_0$ , so agent  $L$  receives the same number of items as agent  $R$ ; and (2) there is exactly one  $R$  move at states  $L_1, \dots, L_{K-1}$ . This proves the lemma, since  $\mathcal{A}^*(K)$  is the only algorithm that satisfies these two conditions. We utilize the fact that the envy of an allocation sequence can be calculated from the number of  $L$  and  $R$  moves in every state: at state  $L_d$ , an  $L$  move increases the envy of agent  $R$  by  $v_d$  while an  $R$  move decreases it by  $v_d$ .

We start with the first property: suppose that the state at time  $T'$  is not 0. Let  $t$  be the last index such that  $s_t = L$ . Allocating  $s_t = R$  instead (and  $s_\ell = R$  for the remaining steps  $\ell > t$ ) reduces the envy of agent  $R$  without entering state  $R_1$ , a contradiction.

For the second property, it suffices to show that if  $s_t = L$  and  $s_{t+1} = R$ , then it must be that at step  $t$  the state is  $L_{K-1}$  (and therefore at step  $t+1$  the state is  $L_K$ ). Assume this is not the case, and we have such a  $t$  where the algorithm is in state  $L_{\hat{K}-1}$ ,  $\hat{K} < K$ . Let  $\ell$  be a step in which

the algorithm is in state  $L_{K-1}$ , which exists by the definition of  $\mathcal{A}(K)$ . Assume that  $\ell > t + 1$  (an analogous argument can be applied to the case that  $\ell < t$ ). We divide  $T'$  into three phases: (1) the first  $t - 1$  items, (2) the next  $\ell - (t + 1)$  items, and (3) the last  $T' - \ell + 2$  items and consider  $s' = s_1, \dots, s_{t-1}, s_{t+2}, \dots, s_\ell, s_t, s_{t+1}, s_{\ell+1}, \dots, s_{T'}$ . Notice that  $s'$  is  $s$ , except the alternating allocations  $L$  then  $R$  are now made at state  $L_{K-1}$  instead of at  $L_{\hat{K}-1}$ . By construction, sequence  $s'$  never goes past state  $L_K$ . We now prove that, using  $s'$ , the envy decreases with respect to  $s$  at each time step after  $t - 1$ , contradicting the assumption  $\mathcal{A}(K) \in \text{OPT}(K)$ .

In phase (1), the envy is unchanged. For phase (2), when using  $\mathcal{A}(K)$ , the pair of moves  $s_t$  and  $s_{t+1}$  increases envy by  $v_{\hat{K}} - v_{\hat{K}-1}$ . Hence, in comparison,  $s'$  has that much less envy during each time step of phase (2). At the start of phase (3) in  $s'$ , the alternating allocations are performed at state  $L_{K-1}$ , increasing envy (in  $s'$ ) by  $v_{K-1} - v_K < v_{\hat{K}} - v_{\hat{K}+1}$ . At all remaining steps in (3), the envy is smaller in  $s'$  (compared to  $s$ ) by  $(v_{\hat{K}} - v_{\hat{K}+1}) - (v_{K-1} - v_K)$ . This completes the proof that  $\mathcal{A}(K)$  must satisfy both properties; the lemma follows.  $\square$

## EC.5. Details Omitted from Section 4.1

### EC.5.1. Proof of Theorem 4

Dickerson et al. (2014) consider a *static* setting with  $T$  items and  $n$  agents where  $v_{it}$  is drawn from a distribution  $D_i$ . They show that, under mild conditions on the distributions, an envy free allocation exists with probability 1 as  $T \rightarrow \infty$  as long as each agent receives roughly  $T/n$  goods, and each agent has higher expected utility for the good they are allocated than those they are not allocated.

**THEOREM EC.1 (Dickerson et al. (2014)).** *Assume that for all  $i, j \in \mathcal{N}$  and items  $t$  the input distributions satisfy (1)  $\Pr[\arg \max_{k \in \mathcal{N}} v_{kt} = \{i\}] = 1/n$ , and (2) there exist constants  $\mu, \mu^*$  such that*

$$0 < \mathbb{E}[v_{it} | \arg \max_{k \in \mathcal{N}} v_{kt} = \{j\}] \leq \mu < \mu^* \leq \mathbb{E}[v_{it} | \arg \max_{k \in \mathcal{N}} v_{kt} = \{i\}].$$

*Then for all  $n \in O(K/\ln K)$ , allocating each item to the agent with the highest value is envy free with probability 1 as  $K \rightarrow \infty$ .*

Given this result, it is straightforward to prove Theorem 4.

**THEOREM 4.** *Algorithm 1 outputs an allocation that is always Pareto efficient. Furthermore, for all  $\varepsilon > 0$ , there exists  $T_0 = T_0(\varepsilon)$ , such that if  $T \geq T_0$ , the output allocation satisfies pairwise EF1 or is envy-free with probability at least  $1 - \varepsilon$ .*

Proof of Theorem 4. Allocating to the agent with the highest value maximizes social welfare and generates a Pareto efficient allocation. When  $D$  is a point mass,  $v_{it} = v$  for all  $i \in \mathcal{N}$  and all  $t \in \mathcal{G}$ , therefore allocating the arriving items in a round-robin manner is EF1. As observed by [Kurokawa et al. \(2016\)](#) in a static setting, proving the following two properties about our allocation algorithm allows the proof of Theorem EC.1 of [Dickerson et al. \(2014\)](#) to go through. The first property is that the probability that agent  $i$  wins item  $t$  is  $1/n$  for all agents  $i$  and items  $t$ . This is satisfied, as all agent values are drawn from the same distribution and tie-breaking is done randomly. The second property is that for some constant  $\Delta > 0$ , and for all agents  $i, j \in \mathcal{N}$  where  $i \neq j$ :  $\mathbb{E}[v_{it} \mid i \text{ receives } t] - \mathbb{E}[v_{it} \mid j \text{ receives } t] \geq \Delta$ .

We now show that when the distribution  $D$  is not a point mass, allocating items to the agent  $i$  with the maximum value  $v_{id}$  and breaking ties uniformly at random ensures that for some constant  $\Delta > 0$ , and for all agents  $i, j \in \mathcal{N}$  where  $i \neq j$ ,  $\mathbb{E}[v_{it} \mid i \text{ receives } t] - \mathbb{E}[v_{it} \mid j \text{ receives } t] \geq \Delta$ .

We largely follow the proof of Lemma 3.2 by [Kurokawa et al. \(2016\)](#) with some simplifications and importantly, handling the case when the distribution  $D$  is discrete.

Since agents' value distributions are identical, we can restate this as:

$$\mathbb{E}[v_{it} \mid i \text{ receives } t] - \mathbb{E}[v_{it} \mid i \text{ does not receive } t] \geq \Delta.$$

Agents are all identical and the algorithm always allocates the item to an agent with the maximum value for it. This implies that  $\mathbb{E}[v_{it} \mid i \text{ receives } t] = \mathbb{E}[\max(v_{1t}, \dots, v_{nt})]$ . Next, if  $D$  is not a point mass, we know that  $\mathbf{Var}[D] > 0$ . From here, we can show that  $\mathbb{E}[\max(v_{1t}, \dots, v_{nt})] > \mathbb{E}[v_{it}]$ .

Let  $\mathbf{Var}[D] = c$ . Let  $\bar{X} = \mathbb{E}[X]$ ,  $p = \Pr[X < \bar{X}]$ . Observe that when  $\mathbf{Var}[D] > 0$ ,  $p \in (0, 1)$ .

$$c = \mathbf{Var}[D]$$

$$\begin{aligned}
&= \mathbb{E}[(X - \bar{X})^2] \\
&\leq \mathbb{E}[|X - \bar{X}|] \\
&= p\mathbb{E}[\bar{X} - X | X < \bar{X}] + (1-p)\mathbb{E}[X - \bar{X} | X \geq \bar{X}] \\
&= -p\mathbb{E}[X | X < \bar{X}] + (1-p)\mathbb{E}[X | X \geq \bar{X}] + (2p-1)\bar{X}
\end{aligned}$$

From here, we analyze the cases where  $p \leq 1/2$  or  $p > 1/2$  separately.

Suppose that  $p > 1/2$ , we use the substitution  $\bar{X} = p\mathbb{E}[X | X < \bar{X}] + (1-p)\mathbb{E}[X | X \geq \bar{X}]$  and then we rewrite the above as:

$$\begin{aligned}
c &\leq 2p(\bar{X} - \mathbb{E}[X | X < \bar{X}]) \\
\implies \frac{c}{2} &\leq \bar{X} - \mathbb{E}[X | X < \bar{X}] \leq \mathbb{E}[X | X \geq \bar{X}] - \mathbb{E}[X | X < \bar{X}]
\end{aligned}$$

Similarly, for  $p \leq 1/2$ ,

$$\begin{aligned}
c &\leq 2(1-p)(\mathbb{E}[X | X \geq \bar{X}] - \bar{X}) \\
\implies \frac{c}{2} &\leq \mathbb{E}[X | X \geq \bar{X}] - \bar{X} \leq \mathbb{E}[X | X \geq \bar{X}] - \mathbb{E}[X | X < \bar{X}]
\end{aligned}$$

Finally,

$$\begin{aligned}
\mathbb{E}[\max(v_{1t}, \dots, v_{nt})] &\geq (1-p^n)\mathbb{E}[X | X \geq \bar{X}] + p^n\mathbb{E}[X | X < \bar{X}] \\
&\geq (1-p)\mathbb{E}[X | X \geq \bar{X}] + p\mathbb{E}[X | X < \bar{X}] \\
&\quad + (p-p^n)(\mathbb{E}[X | X \geq \bar{X}] - \mathbb{E}[X | X < \bar{X}]) \\
&\geq (1-p)\mathbb{E}[X | X \geq \bar{X}] + p\mathbb{E}[X | X < \bar{X}] + (p-p^n)\frac{c}{2} \\
&= \mathbb{E}[v_{it}] + (p-p^n)\frac{c}{2}
\end{aligned}$$

Thus, we have that  $\mathbb{E}[v_{it} | i \text{ receives } t] \geq \mathbb{E}[v_{it}] + (p-p^n)\frac{c}{2}$ . From law of total expectation, we know that  $\mathbb{E}[v_{it}] = \frac{1}{n}\mathbb{E}[v_{it} | i \text{ receives } t] + \frac{n-1}{n}\mathbb{E}[v_{it} | i \text{ does not receive } t]$ . We can combine and rearrange to also show that  $\mathbb{E}[v_{it} | i \text{ does not receive } t] \leq \mathbb{E}[v_{it}] - (p-p^n)\frac{c}{2(n-1)}$ , which allows us to conclude

$$\mathbb{E}[v_{it} | i \text{ receives } t] - \mathbb{E}[v_{it} | i \text{ does not receive } t] \geq (p-p^n) \cdot \frac{c}{2(n-1)} + (p-p^n)\frac{c}{2}$$

Setting  $\Delta$  to  $(p-p^n)\frac{c}{2(n-1)} + (p-p^n)\frac{c}{2}$ , which is positive since  $p \in (0, 1)$  and  $c > 0$ , completes the proof.  $\square$

### EC.5.2. Detail Missing from Proof of Theorem 5

Define

$$c^* = \min \left\{ \min_{i \in \mathcal{N}, j \in [m]: Y_{ij} > Y'_{ij}} \frac{X_{ij}}{Y_{ij} - Y'_{ij}}, \min_{i \in \mathcal{N}, j \in [m]: Y_{ij} < Y'_{ij}} \frac{1 - X_{ij}}{Y'_{ij} - Y_{ij}} \right\}.$$

$Y$  and  $Y'$  are both feasible and strictly different allocations, so there must be some agent  $i$  and item type  $\gamma_j$  such that  $Y_{ij} > Y'_{ij}$  and other some other pair  $\hat{i}, \hat{j}$  such that  $Y_{\hat{i}\hat{j}} > Y'_{\hat{i}\hat{j}}$ . Furthermore, since  $Y_{ij} > 0$  implies that  $X_{ij} > 0$  (recall that an item  $t$  of type  $\gamma_j$  is in  $A_i$  only if  $X_{ij} > 0$ ), we have that  $c^* > 0$ . It remains to show that  $X + c^* \Delta$  is feasible. There are three cases for an agent  $i$  and item type  $\gamma_j$ : (1)  $Y_{ij} = Y'_{ij}$ , (2)  $Y_{ij} > Y'_{ij}$ , (3)  $Y_{ij} < Y'_{ij}$ . For case (1), trivially  $(X + c^* \Delta)_{ij} = X_{ij} \in [0, 1]$ . For case (2),  $(X + c^* \Delta)_{ij} < X_{ij} \leq 1$ , and  $(X + c^* \Delta)_{ij} \geq X_{ij} + \frac{X_{ij}}{Y_{ij} - Y'_{ij}} \cdot (Y'_{ij} - Y_{ij}) = 0$ . Finally, for case (3),  $(X + c^* \Delta)_{ij} > X_{ij} \geq 0$ , and  $(X + c^* \Delta)_{ij} > X_{ij} + \frac{1 - X_{ij}}{Y'_{ij} - Y_{ij}} \cdot (Y'_{ij} - Y_{ij}) \leq 1$ .

## EC.6. Details from Section 4.3 - Beyond Vanishing Envy

### EC.6.1. Example Showing Strict Envy-Freeness and Pareto Efficiency is not Achievable

Here, we give an instance where strict envy-freeness and Pareto efficiency are not achievable even though agents' valuations are not identical up to a multiplicative factor.

**Table EC.1** Instance with three agents.

Item	$g_1$	$g_2$	$g_3$
Value of agent 1	1	1	1
Value of agent 2	0.5	1	1
Value of agent 3	0.25	1	1

In Table EC.1, none of the three agents are identical. However, we claim that in any envy-free and Pareto efficient allocation, agents 2 and 3 will be indifferent towards each other's allocation. Intuitively, the problem is that agents 2 and 3 have identical valuations over the items they could possibly receive in an envy-free and Pareto efficient allocation, items 2 and 3.

Let  $X$  be any Pareto efficient and envy-free allocation. Since it is envy-free and agent 1 is indifferent between the items,  $X_{11} + X_{12} + X_{13} \geq 1$ .

Next, we show that if  $X$  is Pareto efficient, then  $X_{11} = 1$ . If  $X_{11} < 1$ , then either  $X_{12} > 0$  or  $X_{13} > 0$ . In addition, either  $X_{21} > 0$  or  $X_{31} > 0$ . Without loss of generality, suppose  $X_{12} > 0$  and  $X_{21} > 0$ . Then, letting  $c = \min(X_{21}, X_{12})$ , observe that allocation  $X'$  Pareto dominates  $X$ , so  $X$  is not Pareto efficient, where

$$X' = X + \begin{bmatrix} c & -c & 0 \\ -c & c & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

Therefore, we know  $X_{11} = 1$  and as a result,  $X_{21} = X_{31} = 0$ . Finally, because agents 2 and 3 have identical valuations for items 2 and 3, we have  $X_{22} + X_{23} \geq 1$  and  $X_{32} + X_{33} \geq 1$ . Together, this implies that  $X_{22} + X_{23} = X_{32} + X_{33} = 1$  so agents 2 and 3 will be indifferent towards each other's allocations.

### EC.6.2. Proof of Theorem 7

In this section we prove our main structural result.

**THEOREM 7** *Given any instance with  $m$  divisible items and  $n$  additive agents, there always exists an allocation that is simultaneously clique identical strongly envy free (CISEF) and Pareto efficient.*

We prove Theorem 7 by building on a standard approach for finding an envy-free and Pareto efficient allocation, namely solving the Eisenberg-Gale convex program (henceforth E-G program).

Recall that the E-G program with “budgets”  $\mathbf{e}$  is the following.

$$\max \sum_{i=1}^n e_i \log \sum_{j=1}^m v_{ij} x_{ij}, \quad \text{subject to} \quad \sum_{i=1}^n x_{ij} \leq 1, \forall j \in [m], \quad \text{and} \quad x_{ij} \geq 0, \forall i \in [n], j \in [m].$$

We use  $\mathbf{x}$  and  $X$  interchangeably for the allocation. When  $e_i = e_j$  for all  $i, j$ , the outcome is also known as the competitive equilibrium from equal incomes (CEEI), which is envy-free (Varian 1974). There exists a solution to the E-G program with primal variables  $\mathbf{x}$  and dual variables  $\mathbf{p}$  (dual variable  $p_j \geq 0$  corresponds to the first constraint above) satisfying the following conditions.

$$\forall j \in [m] : p_j > 0 \implies \sum_{i=1}^n x_{ij} = 1 \tag{EC.1}$$

$$\forall i \in [n], j \in [m] : \frac{v_{ij}}{p_j} \leq \frac{\sum_{k=1}^m v_{ik}x_{ik}}{e_i} \quad (\text{EC.2})$$

$$\forall i \in [n], j \in [m] : x_{ij} > 0 \implies \frac{v_{ij}}{p_j} = \frac{\sum_{k=1}^m v_{ik}x_{ik}}{e_i} = \max_{k \in [m]} \frac{v_{ik}}{p_k}. \quad (\text{EC.3})$$

These conditions are both necessary and sufficient for a feasible solution to be optimal, and can be derived from the KKT conditions; for completeness we show the derivation in [EC.6.2.1](#). The standard interpretation is that  $e_i$  is the budget of agent  $i$  and  $p_j$  is the price for item  $j$ . Then, a solution  $\mathbf{x}, \mathbf{p}$  consists of prices  $\mathbf{p}$  and allocations  $\mathbf{x}$ , such that each agent spends their entire budget  $e_i$  on “optimal” items and all items are completely sold. We say that an item  $j$  is “optimal” for agent  $i$  given prices  $\mathbf{p}$ , when it maximizes the ratio  $v_{ij}/p_j$ , also known as the *bang-per-buck*.

Given a solution  $\mathbf{x}, \mathbf{p}$ , we write  $X_i$  for the allocation of agent  $i$ , and we say that item  $k$  is allocated to  $i$ ,  $k \in X_i$ , if  $x_{ik} > 0$ . We assume without loss of generality that for any solution,  $\mathbf{x}, \mathbf{p}$ , we have  $\forall j : p_j > 0$ . This holds as long as each item has at least one agent who values it; if this is not the case we can safely drop those items. We prove the following, which immediately implies Pareto efficiency and Theorem 7 since the objective function is monotone.

**THEOREM EC.2.** *There exist budgets  $\mathbf{e}$  and an optimal solution  $(\mathbf{x} = X, \mathbf{p})$  to the E-G convex program with budgets  $\mathbf{e}$ , such that  $X$  is clique identical strongly envy-free.*

We start with an allocation  $X$  that is an optimal solution to the E-G convex program with identical agent budgets  $e_i = 1$  for all  $i$ . Then, at a high level, we break the algorithm into two procedures that jointly alter  $\mathbf{x}, \mathbf{p}$  and  $\mathbf{e}$  such that  $\mathbf{x}, \mathbf{p}$  remains an optimal solution to the convex program with budgets  $\mathbf{e}$ , while preserving envy-freeness, until  $X$  satisfies the desired properties. Specifically, the indifference graph  $I(X)$  will end up being a disjoint set of cliques, such that agents in a clique have identical allocations.

**Optimal Transfers.** Given an allocation, let  $r_i := \frac{v_i(A_i)}{e_i}$  be the bang-per-buck of agent  $i$ . We say that agent  $i$  is indifferent towards any item  $k$  for which  $\frac{v_{ik}}{p_k} = r_i$ . We first give a useful property of solutions  $\mathbf{x}, \mathbf{p}$ .



LEMMA EC.2. Given a solution  $\mathbf{x}, \mathbf{p}$  of an E-G program with budgets  $\mathbf{e}$ , for all agents  $i, j$  such that  $e_i = e_j$ ,  $v_i(X_i) = v_i(X_j)$  if and only if  $\forall k \in X_j : \frac{v_{ik}}{p_k} = r_i$ .

First, we show that  $\sum_{k \in X_i} p_k x_{ik} = e_i$  for any agent  $i$ .

$$\sum_{k \in X_i} p_k x_{ik} \stackrel{\text{KKT condition EC.3}}{=} \sum_{k \in X_i} \frac{v_{ik} \cdot e_i}{\sum_{k' \in X_i} v_{ik'} x_{ik'}} x_{ik} = e_i. \quad (\text{EC.4})$$

For the “only if” direction, we know for all items  $k \in X_j$ ,  $v_{ik} = r_i p_k$ . We can substitute this into our previous equation to get:

$$v_i(X_i) = \sum_{k \in X_i} v_{ik} x_{ik} \stackrel{\text{Cond. EC.3}}{=} \sum_{k \in X_i} r_i p_k x_{ik} = r_i e_i = r_i e_j \stackrel{\text{Eq. EC.4}}{=} \sum_{k \in X_j} r_i p_k x_{jk} = \sum_{k \in X_j} v_{ik} x_{jk} = v_i(X_j).$$

For the “if” direction, assume that there is an item  $k^* \in X_j$  such that  $\frac{v_{ik^*}}{p_{k^*}} \neq r_i$ . Then, KKT condition EC.3 implies that  $v_{ik^*} < r_i p_{k^*}$ , since  $k^*$ ’s bang-per-buck is at most  $r_i$ . Thus, we have:

$$v_i(X_i) = \sum_{k \in X_i} v_{ik} x_{ik} = \sum_{k \in X_i} r_i p_k x_{ik} = r_i e_i = r_i e_j = \sum_{k \in X_j} r_i p_k x_{jk} > \sum_{k \in X_j} v_{ik} x_{jk} = v_i(X_j). \quad \square$$

Lemma EC.2 essentially tells us that under equal budgets, if agent  $i$  is indifferent to agent  $j$ ’s allocation, then  $j$ ’s items are maximum bang-per-buck items for  $i$  as well. This gives some intuition for our approach. Assuming equal budgets, if we move items along indifference edges, we can avoid violating the KKT conditions and our solution will still be optimal. We formalize this idea in Lemma EC.3, but first need the following definition. Given a solution  $(\mathbf{x} = X, \mathbf{p})$  for budgets  $\mathbf{e}$ , we consider a change in allocation of items  $\Delta$ , where  $\Delta_{ik}$  is the change in the allocation of item  $k$  for agent  $i$ .

DEFINITION EC.1 (OPTIMAL TRANSFER). A transfer of items  $\Delta$  is an *optimal transfer* if for all items  $k$ : (1)  $\sum_{i \in [n]} \Delta_{ik} = 0$ , i.e. the total allocation of item  $k$  remains unchanged, (2)  $x_{ik} + \Delta_{ik} \in [0, 1]$  for all agents  $i$ , i.e.  $\mathbf{x} + \Delta$  is feasible, and (3) for all agents  $i$  such that  $\Delta_{ik} > 0$ ,  $\frac{v_{ik}}{p_k} = r_i$ , i.e. if agent  $i$  is given more of item  $k$ , then item  $k$  maximizes bang-per-buck for agent  $i$ .

LEMMA EC.3. Let  $(\mathbf{x} = X, \mathbf{p})$  be a solution for budgets  $\mathbf{e}$  and  $\Delta$  be an optimal transfer. Let  $\delta$  represent the change in budget where  $\delta_i = \sum_{k \in \mathcal{G}} p_k \cdot \Delta_{ik}$ . Let  $X' = X + \Delta$ . Then  $(\mathbf{x}' = X', \mathbf{p})$  is a solution for budgets  $\mathbf{e}' = \mathbf{e} + \delta$ .

Proof of Lemma EC.3.  $(\mathbf{x}, \mathbf{p}, \mathbf{e})$  must satisfy the KKT conditions.  $\mathbf{p}$  does not change and  $\mathbf{x}'$  is feasible by definition; it remains to show that  $(\mathbf{x}', \mathbf{p}, \mathbf{e}')$  satisfies the KKT conditions.

KKT condition EC.1 is satisfied since  $\sum_{i \in [n]} \Delta_{ik} = 0$  for all items  $k$ . Furthermore, notice that  $\Delta_{ik} \neq 0$  implies that  $\frac{v_{ik}}{p_k} = r_i$ : if  $\Delta_{ik} > 0$  this fact is implied by the definition of an optimal transfer, while if  $\Delta_{ik} < 0$ , for  $\mathbf{x} + \Delta$  to be feasible, it must be that  $x_{ik} > 0$ , so  $\frac{v_{ik}}{p_k} = r_i$  is implied by KKT condition EC.3. Thus:

$$\begin{aligned} \frac{\sum_{k=1}^m v_{ik} x'_{ik}}{e'_i} &= \frac{\sum_{k=1}^m v_{ik} x_{ik} + \sum_{k=1}^m v_{ik} \Delta_{ik}}{e_i + \delta_i} = \frac{\sum_{k=1}^m v_{ik} x_{ik} + \sum_{k=1}^m r_i p_k \Delta_{ik}}{e_i + \delta_i} \\ &= \frac{\sum_{k=1}^m v_{ik} x_{ik} + r_i \delta_i}{e_i + \delta_i} = \frac{\sum_{k=1}^m v_{ik} x_{ik}}{e_i}, \end{aligned}$$

where the last equality is implied by the fact that  $r_i = v_i(X_i)/e_i$ . Therefore, the RHS of KKT condition EC.2 does not change (the LHS of course didn't change since it only has values and prices), so KKT condition EC.2 is still satisfied. In addition, when  $x_{ik} > 0$ , KKT condition EC.3 is satisfied by similar reasoning. Finally, it is possible that  $x_{ik} = 0$  but  $x'_{ik} > 0$ . In this case, we know that  $\Delta_{ik} > 0$ , and therefore  $\frac{v_{ik}}{p_k} = r_i = \frac{\sum_{k=1}^m v_{ik} x_{ik}}{e_i}$ , by the definition of an optimal transfer and the definition of  $r_j$ . Thus, KKT condition EC.3 is satisfied.  $\square$

**Indifference Edge Elimination.** For an allocation  $X$  and subset of agents  $S \subseteq \mathcal{N}$ , we overload notation, and let  $X_S$  refer to the allocation for agents in  $S$ . A set of agents  $S$  have identical budgets under  $\mathbf{e}$  if, for all  $i \in S$ ,  $e_i = c$  for some  $c$ . Recall that for an allocation  $X$ ,  $I(X)$  refers to the indifference graph, a graph where we have a vertex for every agent and an edge from (the vertices corresponding to) agent  $i$  to agent  $j$  if  $v_i(X_i) = v_i(X_j)$ . In the remainder of this section we refer to agents and vertices interchangeably. Also, recall that for a directed graph  $G = (V, E)$ , a clique is a subset of vertices  $S \subseteq V$  such that for all  $v \in S, u \in S$ , where  $u \neq v$ , there is an edge

$(u, v) \in E$ , and that a weakly connected component (henceforth just component)  $S$  is a subset of the agents such that for each pair of agents  $i, j \in S$ , there is either a path from  $i$  to  $j$  or a path from  $j$  to  $i$ , and  $S$  is a maximal subgraph with this property.

**DEFINITION EC.2 (CLIQUE ACYCLIC GRAPH).** A directed graph  $G = (V, E)$  is clique acyclic if the vertices can be partitioned into cliques  $C_1, \dots, C_k$ , that is,  $C_i \subseteq V$  for all  $i$ ,  $\cup_{i=1}^k C_i = V$ , and  $C_i \cap C_j = \emptyset$  for all  $i \neq j$ , and where for any cycle  $K$  in the graph,  $K$  only contains vertices from  $C_i$  for some  $i$ .

A crucial step towards producing a CISEF allocation will be to find an allocation  $X$  such that  $I(X)$  is clique acyclic and envy-free, where agents in each clique have the same allocation.

**LEMMA EC.4.** *There exists an algorithm that takes as input a solution  $(\mathbf{x} = X, \mathbf{p})$  for budgets  $\mathbf{e}$  and a component  $S$  with identical budgets, and finds a solution  $(\mathbf{x}' = X', \mathbf{p})$  for budgets  $\mathbf{e}$  where  $I(X'_S)$  is clique acyclic and agents in each clique have identical allocations without violating envy-freeness or adding new indifference edges to  $I(X)$ .*

Take the allocation  $X$  and indifference graph  $I(X_S)$ . At a high level we attempt to apply one of the following two operations; we only apply operation 2 only if operation 1 cannot be done.

- **Operation 1:** Eliminate every cycle that is not a clique.
- **Operation 2:** Partition the graph into cliques by merging cliques and “re-balancing” allocations.

We explain in detail how these operations work, prove they satisfy some basic properties and proceed to use them to prove the Lemma [EC.4](#).

*Operation 1.* Without loss of generality, suppose we have a cycle  $K = (1, \dots, k)$ , where edges go from agent  $i$  to agent  $i + 1$  (modulo  $k$ ). If there exists an  $i$  where there is no edge from  $i - 1$  to  $i + 1$ , we show how to eliminate at least one edge from the cycle while keeping  $v_i(X_i)$  unchanged for all agents and not creating new indifference edges.

To perform this operation notice that the absence of an edge from  $i - 1$  to  $i + 1$ , by Lemma [EC.2](#), implies that there exists an item  $\ell \in X_{i+1}$  that is not a bang-per-buck item for agent  $i - 1$ , i.e.

$\frac{v_{i-1,\ell}}{p_\ell} < r_{i-1}$ . We construct an optimal transfer  $\Delta$ , parameterized by a budget  $b$ , as follows. Agent  $i$  will take  $b$  worth of item  $\ell$  (specifically  $\frac{b}{p_\ell}$  units) from agent  $i+1$ . All other agents  $i' \in K$ ,  $i' \neq i$ , will take arbitrary items of total worth  $b$  from  $i'+1$ . Crucially, for all of the transfers, since there is an  $(i', i'+1)$  edge for all  $i' \in K$ , i.e.  $v_{i'}(X_{i'}) = v_{i'}(X_{i'+1})$ , Lemma EC.2 implies the items taken maximize bang-per-buck for  $i'$ . Next, notice that we can find a  $b > 0$  small enough that  $X + \Delta$  is a feasible allocation. In particular, letting  $e_S$  be the budget of agents in  $S$ , we can choose any  $b \leq \min(e_S, x_{i+1,l} \cdot p_l) = x_{i+1,l} \cdot p_l$  (i.e we just need to ensure agent  $i-1$  does not take more than  $x_{i+1,l}$  of item  $l$ ). Finally, our transfers preserve the total allocation of each item and thus,  $\Delta$  is an optimal transfer. Each agent loses and gains  $b$  worth of items so there is no change in budget associated with  $\Delta$ . Therefore, Lemma EC.3 implies that  $(\mathbf{x}' = X', \mathbf{p})$  is a solution for budgets  $\mathbf{e}$ .

Thus for all agents  $i$ ,  $v_i(X_i) = v_i(X'_i)$ . Next, observe that  $v_{i-1}(X'_i) < v_{i-1}(X_i)$  since it decreases by  $b \cdot r_{i-1}$  but increases by strictly less than  $b \cdot r_{i-1}$ . Since  $v_{i-1}(X_{i-1})$  stays the same, this implies that the indifference edge from agent  $i-1$  to agent  $i$  disappears. Finally, we want to ensure that we do not violate envy-freeness or add new indifference edges. There are two cases. The first case is that  $i$  is indifferent to  $j$  in  $X$ : we only modify allocations in  $S$ , and  $S$  is a component, so we only consider when  $i \in S, j \in S$ . Because  $(\mathbf{x}', \mathbf{p})$  is a solution for budgets  $\mathbf{e}$ , and agents in  $S$  have identical budgets, agent  $i$  will still not envy  $j$  in  $X'$ . Otherwise, the second case is that  $i$  is not indifferent to  $j$  in  $X$  and we want to find a  $b > 0$  small enough such that  $i$  will not envy or be indifferent to  $j$  in  $X'$ . For any pair  $(i, j)$ , we are concerned about  $v_i(A_i) - v_i(A_j)$ . Operation 1 guarantees that  $v_i(X_i) = v_i(X'_i)$ , so  $v_i(A_i) - v_i(A_j)$  can change only when  $j \in K$ . Define  $c = \max_{i \in \mathcal{N}, k \in [m]} \frac{v_{ik}}{p_k}$ : the maximum bang-per-buck for any agent. Then, choose  $b$  such that

$$b < \min \left\{ \frac{v_i(A_i) - v_i(A_j)}{c} : i \in \mathcal{N}, j \in K, v_i(A_i) - v_i(A_j) > 0 \right\}.$$

The budget constraint of  $b$  ensures that  $v_i(A_j)$  will change by at most  $b \cdot c < v_i(A_i) - v_i(A_j)$ .

We can repeatedly use the above process to eliminate all non-clique cycles by eliminating cycles in order of size. Suppose that all cycles of size  $k$  form a clique. It is then possible to ensure that for all cycles of size  $k+1$ , the vertices form a clique. To see why, consider an arbitrary cycle  $K$  of

size  $k + 1$ . If there exists an  $i$  such that there is no edge from agent  $i - 1$  to agent  $i + 1$ , we can eliminate the cycle. Otherwise, for all  $i$ , there is an edge from agent  $i - 1$  to agent  $i + 1$ . Then, we know that for all  $i$ ,  $K \setminus \{i\}$  is a cycle of length  $k$ , and therefore a clique of size  $k$ , implying the vertices of  $K$  form a clique of size  $k + 1$ . Note that any cycle of size 2 immediately forms a clique. Therefore, if we repeatedly choose the smallest size non-clique cycle and eliminate it, we will eventually eliminate all cycles that are not a clique.

*Operation 2.* We construct a set of cliques  $C_1, \dots, C_s$  by starting with each vertex in its own clique and arbitrarily merging cliques if the resulting set of vertices would still form a clique. Suppose we merge to form a clique  $C = \{1, 2, \dots, \ell\}$ . Lemma EC.2 implies that for any agent  $i \in C$ ,  $i$  is indifferent to any item in  $X_C = X_1 + \dots + X_\ell$ , that is  $\frac{v_{iz}}{p_z} = r_i$  for all items  $z \in X_C$ . Thus, we can perform the following re-balance operation to form  $X'$ : for each agent  $i \in C$  we set  $X'_i = \frac{1}{|C|}X_C$  and for all other agents  $i \notin C$ ,  $X'_i = X_i$ .

First, note that this is an optimal transfer, by definition. We want to show that we do not violate envy-freeness or add new indifference edges. There are three cases for the existence of an edge from agent  $i$  to agent  $j$ . If both agents are in  $C$ , nothing changes. If  $i \in C$ ,  $j \notin C$ , nothing changes since  $v_i(X_i)$  is unchanged. If  $i \notin C$ ,  $j \in C$ , if  $i$  was indifferent to all agents in  $C$ , nothing changes. Otherwise,  $i$  is not indifferent to some agent in  $C$ , in which case after the re-balance,  $i$  will lose their indifference edge towards all agents in  $C$ .

*Proof of Lemma EC.4.* We start by applying operation 1, which eliminates every cycle that is not a clique. Now, notice that eliminating all non-clique cycles does *not* imply that we can partition the graph into cliques in a way that is clique acyclic. To see this most clearly, consider a 5 vertex instance, where vertices 1, 2, 3 form a clique, and so do vertices 3, 4, 5. All cycles are also cliques, but the graph is still not clique acyclic (the most “tempting” partition has the issue that cycle (1, 2, 3) contains a vertex that belongs to two cliques). This is where operation 2 comes in.

Once we have eliminated all non-clique cycles, if we are not in a clique acyclic graph with the desired properties, we apply operation 2. It is possible that during the execution of operation 2, an

indifference edge will be eliminated. When this happens, we stop with operation 2 and go back to applying operation 1, and so on. Eventually, this process terminates: neither operation creates any new edges, and furthermore, each time we loop we eliminate at least one edge. The two operations preserve the property that  $(\mathbf{x}' = X', \mathbf{p})$  is a solution for budgets  $\mathbf{e}$ , and that  $X'$  is envy-free, without adding new indifference edges. In addition, re-balancing in operation 2 ensures that agents in each clique have identical allocations. It remains to show that  $I(X')$  is clique acyclic upon termination.

Assume for sake of contradiction that there exists a cycle that includes vertices in two different cliques  $C_1, C_2$ . Therefore there exists some edge from  $C_1$  to  $C_2$  and an edge from  $C_2$  to  $C_1$ . Due to the fact that agents in a clique have identical allocations, this implies that there exists agents  $i_1 \in C_1, i_2 \in C_2$  such that  $i_1$  has edges to all of  $C_2$  and  $i_2$  has edges to all of  $C_1$ . Thus, we can construct a cycle, and thus a clique, containing all agents in  $C_1 \cup C_2$ . Then,  $C_1 \cup C_2$  form a clique, contradicting the fact that no more mergers were possible by operation 2.  $\square$

Once we have an allocation with the properties of Lemma EC.4, it remains to eliminate the edges between cliques, while preserving the property that agents in a clique have the same allocation.

LEMMA EC.5. *There exists an algorithm that takes as input an allocation  $X$  and component  $S$  with identical budgets such that  $I(X_S)$  is clique acyclic (with at least one non-clique edge) and agents in each clique have identical allocations, and finds a set of agent budgets  $\mathbf{e}'$  and solution  $(\mathbf{x}' = X', \mathbf{p})$  for budgets  $\mathbf{e}'$ , where  $I(X'_S)$  consists of  $k > 2$  components  $S_1, S_2, \dots, S_k$ , each component consists of agents with identical budgets,  $I(X'_S)$  has strictly fewer edges than  $I(X_S)$ , and the new allocation does not violate envy-freeness or add new indifference edges.*

Proof of Lemma EC.5. For  $I(X_S)$ , we view each clique as a vertex. Consider the graph  $G$ , where each vertex is a clique  $C_i$  (of  $I(X_S)$ ), and there is an edge between  $C_i$  and  $C_j$  if there exists a  $v_i \in C_i, v_j \in C_j$  where  $(v_i, v_j)$  is an edge in  $I(X_S)$ .  $G$  forms a directed acyclic graph, since  $I(X_S)$  is clique acyclic. Let  $C_s$  be a source vertex in  $G$  and let  $Sk = \{C_1, \dots, C_l\}$  be the sink vertices reachable from  $C_s$ . We know a source and sink vertex exists because we assume there is at least one non-clique edge.

Now we return to  $I(X_S)$ . We will create item transfers  $\Delta$ . We find it intuitive to describe  $\Delta$  via a flow in the following graph. Starting from  $G$ , create a new global source vertex  $s$ , and add all  $(s, v)$  edges for all  $v \in C_s$ . Create a new global sink vertex  $t$ , and add a  $(v, t)$  edge for all  $v \in Sk$ . Finally, we let each edge in the graph have infinite capacity. Next, we find a flow of size  $b$  from  $s$  to  $t$ , with the additional constraints that the flow from  $s$  to each  $v \in C_s$  is the same, and for each  $C \in Sk$ , the flow from each  $v \in C$  to  $t$  is the same. We show such a flow exists by starting with an arbitrary flow of size  $b$  and constructing such a flow. Let  $f_{ij}$  be the flow along edge  $(i, j)$ . For the source clique,  $C_s$ , we can update  $f'_{si} = \frac{1}{|C_s|} \sum_{i \in C_s} f_{si}$  so that  $f'$  now satisfies the equal flow constraint for the source clique. Then, we can always ensure the flow  $f'$  is balanced by updating the flow between vertices in  $C_s$ : For each  $i, j \in C_s, i < j$ , set  $f'_{ij} = f_{ij} + \frac{1}{|C_s|} (f_{sj} - f_{si})$ , where negative flows are added as positive flows in the reverse direction. The flow along all other edges remains the same. To show that  $f'$  is a feasible flow, we show flows are balanced for vertices in the source clique. Observe for an agent  $i \in C_s$ :

$$\sum_{j \in \mathcal{N} \cup \{s\}} f'_{ji} - \sum_{j \in \mathcal{N}} f'_{ij} = \sum_{j \in \mathcal{N}} f_{ji} + \frac{1}{|C_s|} \sum_{j \in C_s} f_{sj} + \frac{1}{|C_s|} \sum_{j \in C_s} (f_{si} - f_{sj}) - \sum_{j \in \mathcal{N}} f_{ij} = \sum_{j \in \mathcal{N} \cup \{s\}} f_{ji} - \sum_{j \in \mathcal{N}} f_{ij}.$$

Therefore, if  $f$  is a valid flow, so is  $f'$ . An analogous procedure can be applied to the sink cliques.

We use this flow to guide the item transfers,  $\Delta$ . For each edge with flow, agent  $i$  will take an arbitrary  $f_{ij}$  worth of items from agent  $j$ . The global sink and source vertices are, obviously, excluded. We choose  $b$  small enough such that  $X + \Delta$  is feasible. We can assume without loss of generality that the sum of flows towards any vertex is at most  $b$ , as any extra flow must be part of a cycle and can be eliminated. Then if  $e_S$  is the budget of agents in  $S$ , we can choose any  $b \leq e_S$ . In addition, since we ensure the total allocation of an item is preserved and we only transfer items along indifference edges between agents of the same budget,  $\Delta$  is an optimal transfer. Therefore we can apply Lemma [EC.3](#).

Since there are no item transfers to  $s$ , each agent  $v \in C_s$  with an incoming flow has the same increased budget. Similarly, for each sink clique  $C$  with a positive flow from  $v \in C$  to  $t$ , each member of  $C$  will have the same decreased budget.

Now, if we take each component in the resulting indifference graph, we claim that each agent in the component will have identical budgets. It is sufficient to show that if  $e'_i \neq e'_j$ , then there will not be an edge from  $v_i$  to  $v_j$ . Since initially everyone in  $S$  had the same budget, and item transfers — as well as “budgets” — go from sinks to sources, if  $e'_i < e'_j$ , there was never an edge from  $v_i$  to  $v_j$  to begin with. Now suppose  $e'_i > e'_j$ ; we have

$$v_i(X_i) = \sum_{k \in X_i} v_{ik} x_{ik} \stackrel{\text{Cond. EC.3}}{=} \sum_{k \in X_i} r_i p_k x_{ik} = r_i e'_i > r_i e'_j = \sum_{k \in X_j} r_i p_k x_{jk} \geq \sum_{k \in X_j} v_{ik} x_{jk} = v_i(X_j),$$

which implies no  $(i, j)$  indifference edge.

Since agents in  $C_s$  have a greater budget than all other remaining agents, there will be at least two components. The final property that we need to prove is that the new allocation does not violate envy-freeness or add new indifference edges. We break into cases. First, if agents  $i$  and  $j$  were indifferent, as we only modify allocations in component  $S$ , we only consider when  $i \in S, j \in S$ . We ensure  $v_i(X_i)$  either stays the same or increases, since we transfer items in the opposite direction of indifference edges. For similar reasons,  $v_i(X_j)$  either stays the same or decreases. On the other hand, if agent  $i$  did not envy  $j$  we can set  $b$  small enough such that  $i$  will not envy  $j$  in  $X'$ . More specifically, choose  $b$  with similarly to before. However, unlike before,  $v_i(A_i) - v_i(A_j)$  can change when either  $i \in S$  or  $j \in S$ , as  $v_i(A_i)$  changes for some  $i \in S$ . Setting

$$b < \min \left\{ \frac{v_i(A_i) - v_i(A_j)}{2c} : i, j \in \mathcal{N}, i \in S \vee j \in S, v_i(A_i) - v_i(A_j) > 0 \right\}$$

as the budget constraint is sufficient as it ensures that  $v_i(A_i)$  and  $v_i(A_j)$  can each change by at most  $\frac{v_i(A_i) - v_i(A_j)}{2}$ .  $\square$

### Putting everything together.

Proof of Theorem EC.2. Our overall algorithm first solves the E-G convex program with budgets  $e$  where  $e_i = 1$  for all  $i$ , to find a solution  $(\mathbf{x} = X, \mathbf{p})$ . We keep track of the set of components  $S$  with identical budgets, where initially  $S = \{\mathcal{N}\}$ . We alternate between applying the algorithm of Lemma EC.4 and Lemma EC.5, henceforth procedure 1 and procedure 2. We start by applying the



former to each  $S \in \mathcal{S}$ . It is possible that after applying procedure 1, edges will be eliminated that result in  $S$  being split into multiple components. In this case, let  $f(S)$  be the set of components formed and update  $\mathcal{S} := \bigcup_{S \in \mathcal{S}} f(S)$ . Note that the clique acyclic and identical allocations properties will still be satisfied by each individual component in  $\mathcal{S}$ . We then apply procedure 2 to each  $S \in \mathcal{S}$ . We perform the same update to  $\mathcal{S}$ , where  $f(S)$  is the set of components with identical budgets found after applying procedure 2. We repeat until applying the two procedures does not decrease the number of edges in the graph. Finally, we perform the re-balance operation (described in procedure 1) on each component (or clique, as we will show).

Both procedures do not add edges and reduce the number of edges. There are at most  $n^2 - n$  initial edges, and thus the algorithm terminates. In addition, both procedures produce allocations and budgets where  $(\mathbf{x}' = A', \mathbf{p})$  is a solution to the E-G program for budgets  $\mathbf{e}'$ .

We now show that  $X$  is CISEF. First,  $X$  is envy-free by construction. Next, consider the final graph  $I(X)$ . We claim that each component is a clique where agents have identical budgets. The component must be clique-acyclic after procedure 1 is applied. Meanwhile, no edges were removed by procedure 2 so there could not have been any edges between cliques. Therefore, the component can only consist of one clique. Next, procedure 2 guarantees that each agent in the clique has the same budget. Lemma EC.2 tells us that agents will have identical valuations, up to a multiplicative factor, for items that any agent in the clique receives. The final re-balance operation doesn't alter any of the above properties.  $\square$

We conclude by discussing the computational complexity of the algorithm. We require an exact solution to the E-G program, which is obtainable in strongly polynomial time (Orlin 2010). The edge-elimination steps happen  $O(n^2)$  times. The only possible issue is the number of bits in the solution  $(\mathbf{x}, \mathbf{p})$  and budgets  $\mathbf{e}$ , as the item transfers described in Lemma EC.4 and EC.5 can both increase the length (in bits) of  $\mathbf{x}$  and  $\mathbf{e}$ . This increase depends primarily on the budget transfers  $b$ . We can always choose  $b$  such that  $b$  is equal to  $(v_i(A_i) - v_i(A_j))/4c$  for some  $i, j$ , where  $c$  is a constant that only depends on the instance and the initial solution to the E-G program. Since

$v_i(A_i), v_i(A_j)$  are linear functions of  $\mathbf{x}$ , their representations are a constant (depending on the  $v_{ij}$  and  $n$ ) larger than the elements of  $\mathbf{x}$ . In addition,  $b$  is an additive constant larger than the bit length of the min difference  $v_i(A_i) - v_i(A_j)$ , so performing the transfer of items under budget  $b$  will also only increase the bit length of the elements of  $\mathbf{x}, \mathbf{e}$  by a constant.

**EC.6.2.1. Deriving KKT conditions** Here, we derive the KKT conditions in order to show that they are both necessary and sufficient conditions for optimal solutions to the Eisenberg-Gale convex program.

We first introduce dual variables  $\mathbf{p}, \mathbf{k}$  for the first and second inequality constraints. From stationarity, we have:

$$\nabla_{\mathbf{x}} \sum_{i=1}^n e_i \log \sum_{j=1}^m v_{ij} x_{ij} = \sum_{j=1}^m p_j \nabla_{\mathbf{x}} \sum_{i=1}^n x_{ij} - 1 + \sum_{i=1}^n \sum_{j=1}^m -k_{ij} \nabla_{\mathbf{x}} x_{ij}$$

For each  $i, j$ , we can take the gradient with respect to  $x_{ij}$  to get:

$$\begin{aligned} e_i \frac{v_{ij}}{\sum_{j=1}^m v_{ij} x_{ij}} &= p_j - k_{ij} \\ \frac{v_{ij}}{p_j} &= \frac{\sum_{j=1}^m v_{ij} x_{ij}}{e_i} \left( 1 - \frac{k_{ij}}{p_j} \right) \end{aligned} \quad (\text{EC.5})$$

The primal and dual feasibility conditions tell us that:  $-x_{ij} \leq 0$ ,  $\sum_{i=1}^n x_{ij} - 1 \leq 0$ ,  $p_j \geq 0$ , and  $k_{ij} \geq 0$ . Finally, complimentary slackness tells us that:

$$x_{ij} > 0 \implies k_{ij} = 0 \text{ and } k_{ij} > 0 \implies x_{ij} = 0 \quad (\text{EC.6})$$

$$p_j > 0 \implies \sum_{i=1}^n x_{ij} = 1 \text{ and } \sum_{i=1}^n x_{ij} < 1 \implies p_j = 0 \quad (\text{EC.7})$$

KKT condition [EC.1](#) follows from Equation [EC.7](#). Meanwhile, we show that KKT conditions [EC.2](#) and [EC.3](#) are equivalent to the stationarity condition plus the first two complementary slackness conditions.

**PROPOSITION EC.1.** *For any  $\mathbf{x}, \mathbf{p}$  where  $p_j > 0$ , KKT conditions [EC.2](#) and [EC.3](#) hold if and only if there exists  $\mathbf{k}$  such that Equations [EC.5](#) and [EC.6](#) hold.*

Proof of Proposition EC.1. Consider any  $\mathbf{x}, \mathbf{p}$ .

We first show the forwards direction. We assume that KKT conditions EC.2 and EC.3 hold and give a  $\mathbf{k}$  such that Equations EC.5 and EC.6 hold. Take any  $i \in [n], j \in [m]$ . We set the value of  $k_{ij}$  depending on whether  $x_{ij} > 0$ . Suppose that  $x_{ij} > 0$ . Then observe that setting  $k_{ij} = 0$  will satisfy both the stationarity and complementary slackness conditions. Otherwise,  $x_{ij} = 0$ . In this case, the slackness conditions trivially hold and we just need to show that there exists a  $k_{ij} \geq 0$  such that Equation EC.5 holds. From KKT condition EC.2, we have:

$$\frac{v_{ij}}{p_j} / \frac{\sum_{j=1}^m v_{ij} x_{ij}}{e_i} \leq 1$$

Letting  $c = \frac{v_{ij}}{p_j} / \frac{\sum_{j=1}^m v_{ij} x_{ij}}{e_i}$ , we solve for  $k_{ij}$  which gives  $k_{ij} = p_j(1 - c)$ , which is non-negative.

Next, we show the reverse direction. Assume that there exists  $\mathbf{k}$  such that Equations EC.5 and EC.6 hold. For any  $i \in [n], j \in [m]$ , because we have that  $p_j > 0$  and  $k_{ij} \geq 0$ , this implies that  $1 - \frac{k_{ij}}{p_j} \leq 1$ . Combine this with Equation EC.5 and we get exactly KKT condition EC.2.

Next, we show that KKT condition EC.3 holds. Assume that  $x_{ij} > 0$ . Equation EC.6 tells us that  $k_{ij} = 0$ . Therefore,  $1 - \frac{k_{ij}}{p_j} = 1$  and applying Equation EC.5 tells us that  $\frac{v_{ij}}{p_j} = \frac{\sum_{j=1}^m v_{ij} x_{ij}}{e_i}$ .  $\square$

## References

- Aleksandrov M, Aziz H, Gaspers S, Walsh T (2015) Online fair division: Analysing a food bank problem. *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2540–2546.
- Alon N, Spencer JH (2000) *The Probabilistic Method* (Wiley), 2nd edition.
- Armony M, Ward AR (2010) Fair dynamic routing in large-scale heterogeneous server systems. *Operations Research* 58(3):624–637.
- Arrow KJ, Intriligator M, eds. (1982) *Handbook of Mathematical Economics* (North-Holland).
- Balseiro S, Lu H, Mirrokni V (2021) Regularized online allocation problems: Fairness and beyond. *International Conference on Machine Learning*, 630–639.
- Bansal N, Jiang H, Singla S, Sinha M (2020) Online vector balancing and geometric discrepancy. *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, 1139–1152.

- Barbanel JB (2005) *The geometry of efficient fair division* (Cambridge University Press).
- Barman S, Bhaskar U, Shah N (2020) Optimal bounds on the price of fairness for indivisible goods. *Web and Internet Economics: 16th International Conference, WINE 2020, Beijing, China, December 7–11, 2020, Proceedings*, 356–369 (Berlin, Heidelberg).
- Batani MH, Chen Y, Ciocan DF, Mirrokni V (2022) Fair resource allocation in a volatile marketplace. *Operations Research* 70(1):288–308.
- Bei X, Lu X, Manurangsi P, Suksompong W (2021) The price of fairness for indivisible goods. *Theory of Computing Systems* 65(7):1069–1093, URL <http://dx.doi.org/10.1007/s00224-021-10039-8>.
- Benadè JG, Kazachkov AM, Psomas CA, Procaccia AD (2018) How to make envy vanish over time. *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 593–610.
- Bernstein S (1946) *The Theory of Probabilities* (Gastehizdat Publishing House).
- Bertsimas D, Farias VF, Trichakis N (2011) The price of fairness. *Operations Research* 59(1):17–31.
- Bertsimas D, Farias VF, Trichakis N (2012) On the efficiency-fairness trade-off. *Management Science* 58(12):2234–2250.
- Bertsimas D, Farias VF, Trichakis N (2013) Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research* 61(1):73–87.
- Bogomolnaia A, Moulin H, Sandomirskiy F (2021) On the fair division of a random object. *Management Science* 68(2):1174–1194.
- Brams SJ, Taylor AD (1996) *Fair Division: From Cake-Cutting to Dispute Resolution* (Cambridge University Press).
- Budish E (2011) The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6):1061–1103.
- Butler M, Williams HP (2002) Fairness versus efficiency in charging for the use of common facilities. *Journal of the Operational Research Society* 53(12):1324–1329.
- Caragiannis I, Kaklamanis C, Kanellopoulos P, Kyropoulou M (2009) The efficiency of fair division. *Proceedings of the 5th Conference on Web and Internet Economics (WINE)*, 475–482.

- Caragiannis I, Kurokawa D, Moulin H, Procaccia AD, Shah N, Wang J (2016) The unreasonable fairness of maximum Nash product. *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, 305–322.
- Chaudhury BR, Cheung YK, Garg J, Garg N, Hoefer M, Mehlhorn K (2018) On fair division for indivisible items. *38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2018)* (Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik).
- Cole R, Gkatzelis V (2015) Approximating the Nash social welfare with indivisible items. *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC)*, 371–380.
- Devanur NR, Papadimitriou CH, Saberi A, Vazirani VV (2008) Market equilibrium via a primal–dual algorithm for a convex program. *Journal of the ACM (JACM)* 55(5):1–18.
- Dickerson JP, Goldman JR, Karp J, Procaccia AD, Sandholm T (2014) The computational rise and fall of fairness. *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, 1405–1411.
- Eisenberg E, Gale D (1959) Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics* 30(1):165–168.
- Foley D (1967) Resource allocation and the public sector. *Yale Economics Essays* 7:45–98.
- Freeman R, Zahedi SM, Conitzer V, Lee BC (2018) Dynamic proportional sharing: A game-theoretic approach. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 2(1):3.
- Friedman E, Psomas CA, Vardi S (2015) Dynamic fair division with minimal disruptions. *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, 697–713.
- Friedman E, Psomas CA, Vardi S (2017) Controlled dynamic fair division. *Proceedings of the 2017 ACM Conference on Economics and Computation*, 461–478.
- Gal Y, Mash M, Procaccia AD, Zick Y (2017) Which is the fairest (rent division) of them all? *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 4841–4843.
- Garg J, Hoefer M, Mehlhorn K (2018) Approximating the nash social welfare with budget-additive valuations. *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2326–2340.
- Garg J, Végh LA (2019) A strongly polynomial algorithm for linear exchange markets. *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 54–65.

- Ghodsi A, Zaharia M, Hindman B, Konwinski A, Shenker S, Stoica I (2011) Dominant resource fairness: Fair allocation of multiple resource types. *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation (NSDI)*, 24–37.
- Gkatzelis V, Psomas A, Tan X (2021) Fair and efficient online allocations with normalized valuations. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 5440–5447.
- He J, Procaccia AD, Psomas C, Zeng D (2019) Achieving a fairer future by changing the past. *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 343–349.
- Hoeffding W (1963) Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* 58(301):13–30.
- Kash I, Procaccia AD, Shah N (2014) No agent left behind: Dynamic fair division of multiple resources. *Journal of Artificial Intelligence Research* 51:579–603.
- Kurokawa D, Procaccia AD, Wang J (2016) When can the maximin share guarantee be guaranteed? *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, 523–529.
- Lee MK, Kusbit D, Kahng A, Kim JT, Yuan X, Chan A, See D, Noothigattu R, Lee S, Psomas A, Procaccia AD (2019) WeBuildAI: Participatory framework for algorithmic governance. *Proceedings of the ACM on Human-Computer Interaction* 3:1–35.
- Li B, Li W, Li Y (2018) Dynamic fair division problem with general valuations. *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 375–381.
- Lipton RJ, Markakis E, Mossel E, Saberi A (2004) On approximately fair allocations of indivisible goods. *Proceedings of the 6th ACM Conference on Economics and Computation (EC)*, 125–131.
- Mehta A, Saberi A, Vazirani U, Vazirani V (2007) Adwords and generalized online matching. *Journal of the ACM (JACM)* 54(5):22–es.
- Miller HE, Pierskalla WP, Rath GJ (1976) Nurse scheduling using mathematical programming. *Operations Research* 24(5):857–870.
- Narayan VV, Suzuki M, Vetta A (2021) Two birds with one stone: Fairness and welfare via transfers. Caragiannis I, Hansen KA, eds., *Algorithmic Game Theory*, 376–390.

- Orlin JB (2010) Improved algorithms for computing Fisher’s market clearing prices. *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*, 291–300.
- Procaccia AD (2016) Cake cutting algorithms. Brandt F, Conitzer V, Endress U, Lang J, Procaccia AD, eds., *Handbook of Computational Social Choice*, chapter 13 (Cambridge University Press).
- Ruhe G, Fruhwirth B (1990)  $\epsilon$ -optimality for bicriteria programs and its application to minimum cost flows. *Computing* 44(1):21–34.
- Sanders P (1996) On the competitive analysis of randomized static load balancing. *Proceedings of the 1st Workshop on Randomized Parallel Algorithms (RANDOM)*.
- Su FE (1999) Rental harmony: Sperner’s lemma in fair division. *American Mathematical Monthly* 106(10):930–942.
- Su X, Zenios SA (2006) Recipient choice can address the efficiency-equity trade-off in kidney transplantation: A mechanism design model. *Management Science* 52(11):1647–1660.
- Vardi S, Psomas A, Friedman E (2021) Dynamic fair resource division. *Mathematics of Operations Research* In press.
- Varian H (1974) Equity, envy and efficiency. *Journal of Economic Theory* 9:63–91.
- Walsh T (2011) Online cake cutting. *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*, 292–305.