

Crosscutting Areas

Fair and Efficient Online Allocations

Gerdus Benadè,^{a,*} Aleksandr M. Kazachkov,^b Ariel D. Procaccia,^c Alexandros Psomas,^d David Zeng^e

^aQuestrom School of Business, Boston University, Boston, Massachusetts 02215; ^bDepartment of Industrial & Systems Engineering, University of Florida, Gainesville, Florida 32611; ^cSchool of Engineering and Applied Sciences, Harvard University, Boston, Massachusetts 02134;

^dDepartment of Computer Science, Purdue University, West Lafayette, Indiana 47907; ^eJane Street Capital, New York, New York 10281

*Corresponding author

Contact: benade@bu.edu,  <https://orcid.org/0000-0002-0024-0026> (GB); akazachkov@ufl.edu,  <https://orcid.org/0000-0002-4949-9565>

(AMK); arielpro@seas.harvard.edu,  <https://orcid.org/0000-0002-8774-5827> (ADP); apsomas@cs.purdue.edu,

 <https://orcid.org/0000-0002-7709-5058> (AP); davidzengx@gmail.com (DZ)

Received: June 27, 2022

Revised: April 27, 2023

Accepted: September 25, 2023

Published Online in Articles in Advance:
November 3, 2023

Area of Review: Market Analytics and
Revenue Management

<https://doi.org/10.1287/opre.2022.0332>

Copyright: © 2023 INFORMS

Abstract. 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 nontrivial fairness and nontrivial 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).

Funding: A. Psomas is supported in part by the National Science Foundation [Award CCF-2144208] and Google [the AI for Social Good Award and the Research Scholar Award]. This work was partially supported by the National Science Foundation [Grants IIS-2147187, IIS-2229881, and CCF-2007080] and the Office of Naval Research [Grant N00014-20-1-2488].

Supplemental Material: The e-companion is available at <https://doi.org/10.1287/opre.2022.0332>.

Keywords: online allocation • fair division • fairness-efficiency trade-off

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, Balseiro et al. 2021, Bateni et al. 2022), organ transplantation (Su and Zenios 2006, Bertsimas et al. 2013), nurse shift 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. 2022).

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 that 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, whereas the integral solution that 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 that is EF1 and Pareto efficient for any number of items (Caragiannis et al. 2016), but this requires up-front 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 trade-off 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 the time of allocation and are generated according to one of four different adversary models, which we describe. 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 Sections 3.1 and 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, *nonadaptive* 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 nontrivial efficiency guarantees or nontrivial fairness properties.

In Section 4, we study weaker, Bayesian adversaries. Section 4.1 considers the weakest of these, which 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 that 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 nonidentical, 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 $v_{it'}$ but not with $v_{it''}$. 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 (henceforth, the E-G 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 that 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 in the e-companion).

1.2. Related Work

Our paper is related to the growing literature on *online* or *dynamic* fair division (Walsh 2011; Kash et al. 2014; Aleksandrov et al. 2015; Friedman et al. 2015, 2017; Freeman et al. 2018; Li et al. 2018; He et al. 2019; 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 (i.i.d.). 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 trade-off 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. Gao et al. (2021) study an online version of a Fisher market in which items arrive over time. They define an online equilibrium to be such that the time-averaged prices and allocations form an equilibrium for the corresponding offline market with item supplies proportional to the item arrival probabilities, and they obtain asymptotic fairness guarantees. Similarly, our results for correlated distributions (Theorems 6 and 8) leverage a connection between the online instance and an offline instance in which item type values are scaled by their frequency.

Beyond envy, the *price of fairness* measures the relative loss in social welfare that results from enforcing a fairness constraint. The price of fairness has been studied in static settings for divisible items (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). Our work is similar in spirit; 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 Section EC.1 in the e-companion for a reminder.

We now discuss the different adversary models that 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) and (2) as *worst case* and adversaries (3) and (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. Nonadaptive 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_{it} . For simplicity, we treat this setting as follows. Each item t has one of m types. Agent i has value $v_i(\gamma)$ for an item of type γ ; the type of each item is drawn i.i.d. from a distribution D with support G_D , $|G_D| = m$. We write $f_D(\gamma)$ for the probability that the t th item has type γ . 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 toward j is $\text{Envy}_{i,j}(A) = \max\{v_i(A_j) - v_i(A_i), 0\}$, whereas $\text{Envy}(A) = \max_{i,j \in \mathcal{N}} \text{Envy}_{i,j}(A)$ is the maximum envy. $\text{Envy}(A) = 0$ implies that the allocation is envy free. An allocation A is 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 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 $u > u'$, 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 $u' > u$. Where appropriate, we use a notion of approximate Pareto efficiency, initially by Ruhe and Fruhwirth (1990), to measure the efficiency of our algorithms. An allocation with utility profile u is α -Pareto efficient (for $0 < \alpha \leq 1$) when u/α is undominated.

Because 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 α -Pareto efficient ex post if it always outputs an α -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 α -Pareto efficient ex ante if the expected utility profile is α -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, whereas 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 $\text{Envy}_T = T/2$.

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 that it, too, fails to achieve vanishing envy.

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	$1/2 - \epsilon$	$3/2 - \epsilon$	$3/2 - 2\epsilon$	$5/2 - 2\epsilon$...

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 that 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 ϵ 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 one every two steps. Therefore, the maximum envy at step $2t$ is approximately t , and Envy_T/T approaches $1/2$ as T goes to infinity.

These examples suggest that it is nontrivial to come up with an allocation algorithm that achieves vanishing envy. Coupling vanishing envy with Pareto efficiency, 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 as follow. Is this optimal, or are there other strategies with even stronger fairness properties? We provide an adaptive adversary strategy that 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 nonadaptive adversary, no algorithm can achieve vanishing envy while being $(\frac{1}{n} + \epsilon)$ -Pareto efficient for any $\epsilon > 0$. This clearly establishes the boundaries of what is possible against worst-case adversaries; any allocation algorithm must choose between achieving either nontrivial fairness guarantees or nontrivial 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 $v_{it} = 1$ for all $i \in [n], t \in [T]$. This optimal adversary strategy is nonadaptive, and therefore, because 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, give 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 and 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 1^n . All omitted proofs appear in the e-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 because 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 (decreases) the maximum envy, the adversary benefits by increasing (decreasing) the corresponding valuation to one (to zero).

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 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 nonadaptive 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 Section EC.3 in the e-companion) to Envy_T^{ij} , which equals $\sum_{t=1}^T X_t^{ij}$ when envy exists. It follows that, for $\lambda > 0$,

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

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

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

where the last inequality uses the assumption that $T \geq n \log T$. Because 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 (R), which we will refer to as making an L (R) step, the adversary moves one state to the left (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 $\Omega(T^{r/2})$.

Figure 1. Adversary Strategy for the Two-Agent Lower Bound

Notes. In state L_d , an item valued $(1, v_d)$ arrives, whereas 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.

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.

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 that 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, equal numbers of items are allocated to both agents. All of these items have value 1 to agent R , yielding a net effect of zero 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

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

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 bound is positive for $K \leq \sqrt{T'/2}$, we find the critical point:

$$\begin{aligned} \frac{\partial}{\partial K} (K^r + T' r(1-r)K^{r-2}) &= rK^{r-1} - T' r(1-r)(2-r)K^{r-3} \\ &= 0 \Rightarrow K = \sqrt{T'(1-r)(2-r)}. \end{aligned}$$

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

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

We now show how to extend this adversarial instance to n agents.

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 who value every item at zero. Allocating to agents L or R advances the state of the adversary as before; allocating to an agent $i \in \mathcal{N} \setminus \{L, R\}$ does not affect the state.

Let T_0 be the number of items allocated to one of agents 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 $i \in [3, \dots, n]$ who is allocated $\Omega(T/n)$ items. Without loss of generality, at least half these items were allocated while the adversary was in the left states. This implies that

agent L values each of these items at one, so agent L has total value $\Omega(T/n)$ 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 toward agent i is at least $\Theta(T/n) - o(T/n) \in \Theta(T/n)$. \square

3.3. Nontrivial 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 nonadaptive worst-case adversary, which immediately establishes the result against adaptive adversaries.

Theorem 3. *Against a nonadaptive 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 ε 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 $v_{ij'} = \varepsilon$, so 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 $\text{Envy}(A^t) \leq f(T)$ because 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 $x_i = \frac{f(T)}{\varepsilon} (1 + \frac{2}{\varepsilon})^{i-1}$ bounds the largest deviation from the mean number of items (T/n^2) permissible in segment i subject to the allocation having sublinear envy.

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 $\varepsilon \cdot y$. However, $\text{Envy}_{\hat{k},k}(A^{T/n}) \leq f(T)$, 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 ,

$$v_{\hat{k}}(A_k^{T(i-1)}) - v_{\hat{k}}(A_{\hat{k}}^{T(i-1)}) \geq -\sum_{i'=1}^{i-1} 2x_{i'},$$

where the sum is over the maximum deviations from T/n^2 in previous segments. The bound is tight when \hat{k} received $\frac{T}{n^2} + x_{i'}$ items from each previous segment i' , 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,

$$\begin{aligned} f(T) &\geq \text{Envy}_{\hat{k},k}(A^{T_i}) \geq \varepsilon \cdot y + v_{\hat{k}}(A_k^{T(i-1)}) - v_{\hat{k}}(A_{\hat{k}}^{T(i-1)}) \\ &\geq \varepsilon \cdot y - 2 \sum_{i' < i} x_{i'}, \end{aligned}$$

which after substituting each prior $x_{i'}$ with the bound from the induction hypothesis, 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, because $x_i < x_n$, 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 that A^T is not $(\frac{1}{n} + \varepsilon)$ -Pareto efficient, a contradiction. \square

We use this result to prove Theorem 3 for a nonadaptive adversary.

Proof of Theorem 3. Suppose that there is an allocation algorithm that 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, and the remaining items have no value. We bound the number of items the algorithm can allocate to each agent in each segment by induction; this time our bounds are looser and probabilistic. Let $\mathcal{E}(x_1, \dots, x_{i-1}, x)$ be the event that every agent receives a number of items in $(\frac{T}{n^2} \pm x_j)$ from each segment $j = 1, \dots, i - 1$, and there exists an agent who receives a number of items at distance at least x from $\frac{T}{n^2}$ in segment i .

Let α^* be a number such that $\alpha^* \cdot f(T) \in o(T)$ and $\alpha^* \in \omega(1)$. For example, one may think of $\alpha^* = T^\delta$ for some small $\delta > 0$ that depends on $f(T)$. We show by induction that if $x_i^* \in \sup_x \{Pr[\mathcal{E}(x_1^*, \dots, x_{i-1}^*, x)] \geq \frac{1}{\alpha^*}\}$, then $x_i^* \leq \frac{\alpha^* f(T)}{\epsilon} (1 + \frac{2}{\epsilon})^{i-1}$.

As base case when $i = 1$, consider the allocation of the first segment when the algorithm is faced with instance I_1 . Suppose, from items 1 through $\frac{T}{n}$, the algorithm allocates to some agent k at least $\frac{T}{n^2} + x_1^*$ items, for $x_1^* > 0$, with probability $1/\alpha^*$. The conditional expected envy of some agent \hat{k} (who received fewer than $\frac{T}{n^2}$ items) under $\mathcal{E}(x_1^*)$ is at least $\epsilon \cdot x_1^*$, and $\mathbb{E}[\text{Envy}(A^T)] \geq \frac{1}{\alpha^*} \cdot \epsilon x_1^*$. Because $\mathbb{E}[\text{Envy}(A^T)] \leq f(T)$, we have that $x_1^* \leq \frac{\alpha^* f(T)}{\epsilon} = \frac{\alpha^* f(T)}{\epsilon} (1 + \frac{2}{\epsilon})^0$. The same bound is obtained for deviations below $\frac{T}{n^2}$. Because the first $\frac{T}{n}$ items are identical for all instances, the bound on the number of items received from the segment also holds for instances I_2, \dots, I_n .

Suppose that for all $j = 1, \dots, i - 1$, $x_j^* \leq \frac{\alpha^* f(T)}{\epsilon} (1 + \frac{2}{\epsilon})^{j-1}$. We analyze the envy of the algorithm at the end of I_i under the event $\mathcal{E}(x_1^*, \dots, x_{i-1}^*, x_i^*)$. For similar reasons as before,

$$\begin{aligned} f(T) &\geq \mathbb{E}[\text{Envy}(A^T)] \geq \frac{1}{\alpha^*} \left[x_i^* \cdot \epsilon - \sum_{j=1}^{i-1} 2x_j^* \right] \\ &\geq \frac{1}{\alpha^*} \left[x_i^* \cdot \epsilon - \frac{2\alpha^* f(T)}{\epsilon} \cdot \sum_{j=1}^{i-1} \left(1 + \frac{2}{\epsilon}\right)^{j-1} \right] \\ &\geq \frac{1}{\alpha^*} \left[x_i^* \cdot \epsilon - \alpha^* f(T) \left(\left(1 + \frac{2}{\epsilon}\right)^{i-1} - 1 \right) \right]. \end{aligned}$$

It follows that $x_i^* \leq \frac{\alpha^* f(T)}{\epsilon} (1 + \frac{2}{\epsilon})^{i-1}$, which holds on instances I_1, \dots, I_n , to complete the induction.

Finally, we analyze the efficiency of the algorithm on instance I_n . For arbitrary agent i , let \tilde{v}_{ij} be the value that i has for each item in segment j . We bound their expected utility as

$$u_i \leq \sum_{j=1}^n \left(\frac{T}{n^2} + x_j^* \right) \tilde{v}_{ij} + \frac{n}{\alpha^*} \sum_{k=1}^T v_{ik},$$

where the first term assumes a deviation of at most x_j^* in each segment and the second accounts for the worst-case large deviation in which a single agent receives all items. It now follows that

$$\begin{aligned} u_i &\leq \sum_{j=1}^n \left(\frac{T}{n^2} + x_n^* \right) \tilde{v}_{ij} + \frac{n}{\alpha^*} \left(\frac{T}{n} + \frac{T(n-1)\epsilon}{n} \right), \quad (x_j^* < x_n^* \forall j < n) \\ &= \left(\frac{T}{n^2} + x_n^* \right) \sum_{j=1}^n \tilde{v}_{ij} + \frac{nT}{\alpha^*} \left(\frac{1}{n} + \frac{(n-1)\epsilon}{n} \right), \\ &= \left(\frac{T}{n^2} + x_n^* \right) (1 + (n-1)\epsilon) + \frac{T}{\alpha^*} (1 + (n-1)\epsilon), \\ &= (1 + (n-1)\epsilon) \left(\frac{T}{n^2} + x_n^* + \frac{T}{\alpha^*} \right) \\ &= \frac{T}{n^2} (1 + (n-1)\epsilon) \left(1 + \frac{n^2 x_n^* + n^2}{T} \right) \\ &\leq \frac{T}{n} \left(\frac{1}{n} + \epsilon - \frac{\epsilon}{n} \right) \left(1 + \frac{n^2 \alpha^* f(T)}{\epsilon T} \left(1 + \frac{2}{\epsilon} \right)^{n-1} + \frac{n^2}{\alpha^*} \right). \quad (\text{ind. hyp.}) \end{aligned}$$

By construction, $\alpha^* \in \omega(1)$ and $\alpha^* \cdot f(T) \in o(T)$, from which it follows that $\alpha^* > \frac{2n^3}{\epsilon}$, and eventually, T is large enough to satisfy $\frac{\alpha^* f(T)}{T} < \frac{\epsilon^2}{2n^3(1+\frac{2}{\epsilon})^{n-1}}$. Together, this yields

$$\begin{aligned} u_i &< \frac{T}{n} \left(\frac{1}{n} + \epsilon - \frac{\epsilon}{n} \right) \left(1 + \frac{\epsilon}{2n} + \frac{\epsilon}{2n} \right) \\ &= \frac{T}{n} \left(\frac{1}{n} + \epsilon - \frac{\epsilon}{n} \right) \left(1 + \frac{\epsilon}{n} \right) < \frac{T}{n} \left(\frac{1}{n} + \epsilon \right), \end{aligned}$$

whereas the allocation that gives items $[\frac{T}{n}(i-1) + 1, \dots, \frac{T}{n}i]$ to agent i results in utility $u_i' = \frac{T}{n}$ to each $i \in \mathcal{N}$. 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 nontrivial 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 i.i.d. items. Using a result by Dickerson et al. (2014), we show 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 i.i.d. 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 always outputs an allocation that is Pareto efficient. Furthermore, for every distribution D , at least one of the properties hold.*

1. The output allocation is EF1 for all $T \geq 0$.
2. For all $\varepsilon > 0$, there exists $T_0 = T_0(\varepsilon)$, such that for all $T \geq T_0$, the output allocation is envy free with probability at least $1 - \varepsilon$.

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 in the e-companion.

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 γ_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 γ_j , G_D is the support of D , $|G_D| = m$, and $v_i(\gamma_j)$ is the value of an item of type γ_j to agent i . For ease of notation, we sometimes refer to item type γ_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 $m = |G_D|$ is 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 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 γ_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 $A = (A_1, \dots, A_n)$ be any allocation of S where an item of type γ_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 γ_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 $fr(\gamma_j) = \sum_{t \in S} \mathbb{1}\{\text{item } t \text{ has type } \gamma_j\}$. We prove the theorem in two steps. First, we show that X is Pareto efficient under \tilde{v} . Second, we show that this implies that 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 number 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 Δ' , where $\Delta'_{ij} = \Delta_{ij} \cdot fr(\gamma_j) / f_D(\gamma_j)$. Observe that the change in utilities induced by transfers Δ' under v' equals the change in utilities induced by transfers Δ 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 that c is well defined and that $(X + c\Delta')_{ij} = \delta_j X'_{ij} + (1 - \delta_j) X_{ij} \in [0, 1]$, where $\delta_j = c \cdot fr(\gamma_j) / f_D(\gamma_j) \leq 1$. We conclude that $X + c\Delta'$ 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 γ_j given to agent i in A . Define Y' similarly for A' .

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

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

which is 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 in the e-companion, 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 that 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 γ_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

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

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 $\varepsilon = \frac{1}{n^2 T}$, 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

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

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*, 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 in the e-companion 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 that 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, although 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 CISEF and Pareto efficient.*

This result is constructive and somewhat technical (see Section EC.6.2 in the e-companion). We provide a sketch.

Proof Sketch of Theorem 7. We start by solving 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”

e consists of

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

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 let p and e be the prices and budgets.

Then, at a high level, we proceed by repeatedly altering X, p , and e in such a way that X, p remains an optimal solution to the convex program with budgets 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, Chaudhury et al. 2018, Garg 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 (in Lemma EC.2 in the e-companion) that 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 that 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 γ_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 in the e-companion. Notice in Algorithm 3,

step 3 that there is a unique agent with smallest value (up to tiebreaking) because all agents agree on the value of all items that have gone to the clique (up to multiplicative factors).

Theorem 8. *Algorithm 3 always outputs an ex post Pareto-efficient allocation. Furthermore, for all distributions D and every pair of agents i, j , at least one of the following holds:*

1. *i envies j by at most one item, or*
2. *for all $\varepsilon > 0$, there exists $T_0 = T_0(\varepsilon)$, such that i does not envy j with probability at least $1 - \varepsilon$ when $T \geq T_0$.*

Proof of Theorem 8. Let X^* be the fractional CISEF and Pareto-efficient allocation, and let A be 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., $S = \cup_{t \in C_k} A_t$). 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 the $1/|C_p|$ 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 one 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 $\mathbb{E}[v_i(\tilde{A}_i) - v_i(\tilde{A}_j)] = T\mathbb{E}[v_i(X_i^*) - v_i(X_j^*)] = cT$.

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

5.1. Additivity Assumption

We assume that agents have additive valuations. This common assumption is often considered strong. However, for

the purpose of defining envy in our online setting, we believe that it is quite natural. Because 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 $\text{Envy}_{ij}(A) = \sum_{t \in A_i} v_{it} - \sum_{t \in A_j} v_{it} = \sum_{t=1}^T v_{it}(\mathbb{1}_{t \in A_i} - \mathbb{1}_{t \in A_j})$ as the sum of per-round envies, so assuming additive valuations amounts to assuming that envy is additive over time.

5.2. Computational Considerations

Theorem 7 ensures that 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, p) and budgets e , as the item transfers in Lemmas EC.4 and EC.5 in the e-companion can both increase the length (in bits) of X and e . However, as we discuss in Section EC.6.2 in the e-companion, this increase is limited to a constant number of bits.

5.3. 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 nonadaptive 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 superconstant 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.

References

- Aleksandrov M, Aziz H, Gaspers S, Walsh T (2015) Online fair division: Analysing a food bank problem. Yang Q, Wooldridge M, eds. *Proc. 24th Internat. Joint Conf. Artificial Intelligence (IJCAI)* (AAAI Press, Palo Alto, CA), 2540–2546.
- Alon N, Spencer JH (2000) *The Probabilistic Method*, 2nd ed. (Wiley, New York).
- Armony M, Ward AR (2010) Fair dynamic routing in large-scale heterogeneous server systems. *Oper. Res.* 58(3):624–637.
- Arrow KJ, Intriligator M, eds. (1982) *Handbook of Mathematical Economics* (North-Holland, Amsterdam).
- Balseiro S, Lu H, Mirrokni V (2021) Regularized online allocation problems: Fairness and beyond. Meila M, Zhang T, eds. *Internat. Conf. Machine Learn.* (PMLR, Cambridge, MA), 630–639.
- Bansal N, Jiang H, Singla S, Sinha M (2020) Online vector balancing and geometric discrepancy. Makarychev K, Makarychev Y, Tulsiani M, Kamath G, Chuzhoy J, eds. *Proc. 52nd Annual ACM SIGACT Sympos. Theory Comput.* (ACM, New York), 1139–1152.
- Barbanel JB (2005) *The Geometry of Efficient Fair Division* (Cambridge University Press, Cambridge, UK).

- Barman S, Bhaskar U, Shah N (2020) Optimal bounds on the price of fairness for indivisible goods. Chen X, Gravin N, Hoefer M, Mehta R, eds. *Proc. Web Internet Econom. 16th Internat. Conf. WINE 2020* (Springer, Cham, Switzerland), 356–369.
- Bateni MH, Chen Y, Ciocan DF, Mirrokni V (2022) Fair resource allocation in a volatile marketplace. *Oper. Res.* 70(1):288–308.
- Bei X, Lu X, Manurangsi P, Suksompong W (2021) The price of fairness for indivisible goods. *Theory Comput. Systems* 65(7):1069–1093.
- Benadè JG, Kazachkov AM, Psomas CA, Procaccia AD (2018) How to make envy vanish over time. Elkind E, Vohra R, Tardos E, eds. *Proc. 19th ACM Conf. Econom. Comput. (EC)* (ACM, New York), 593–610.
- Bernstein S (1946) *The Theory of Probabilities* (Gastehizdat Publishing House, Moscow).
- Bertsimas D, Farias VF, Trichakis N (2011) The price of fairness. *Oper. Res.* 59(1):17–31.
- Bertsimas D, Farias VF, Trichakis N (2012) On the efficiency-fairness trade-off. *Management Sci.* 58(12):2234–2250.
- Bertsimas D, Farias VF, Trichakis N (2013) Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Oper. Res.* 61(1):73–87.
- Bogomolnaia A, Moulin H, Sandomirskiy F (2021) On the fair division of a random object. *Management Sci.* 68(2):1174–1194.
- Brams SJ, Taylor AD (1996) *Fair Division: From Cake-Cutting to Dispute Resolution* (Cambridge University Press, Cambridge, UK).
- Budish E (2011) The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *J. Political Econom.* 119(6):1061–1103.
- Butler M, Williams HP (2002) Fairness vs. efficiency in charging for the use of common facilities. *J. Oper. Res. Soc.* 53(12):1324–1329.
- Caragiannis I, Kaklamanis C, Kanellopoulos P, Kyropoulou M (2009) The efficiency of fair division. Leonardi S, ed. *Proc. 5th Conf. Web Internet Econom. (WINE)* (Springer, Berlin, Heidelberg, Germany), 475–482.
- Caragiannis I, Kurokawa D, Moulin H, Procaccia AD, Shah N, Wang J (2016) Conitzer V, Bergemann D, Chen Y, eds. The unreasonable fairness of maximum Nash product. *Proc. 17th ACM Conf. Econom. Comput. (EC)* (ACM, New York), 305–322.
- Chaudhury BR, Cheung YK, Garg J, Garg N, Hoefer M, Mehlhorn K (2018) On fair division for indivisible items. Ganguly S, Pandya P, eds. *38th IARCS Annual Conf. Foundations Software Tech. Theoret. Comput. Sci. (FSTTCS 2018)* (Schloss Dagstuhl, Wadern, Germany), 25:1–25:17.
- Cole R, Gkatzelis V (2015) Approximating the Nash social welfare with indivisible items. Servedio R, Rubinfeld R, eds. *Proc. 47th Annual ACM Sympos. Theory Comput. (STOC)* (ACM, New York), 371–380.
- Devanur NR, Papadimitriou CH, Saberi A, Vazirani VV (2008) Market equilibrium via a primal-dual algorithm for a convex program. *J. ACM* 55(5):1–18.
- Dickerson JP, Goldman JR, Karp J, Procaccia AD, Sandholm T (2014) The computational rise and fall of fairness. Brodley CE, Stone P, eds. *Proc. Twenty-Eighth AAAI Conf. Artificial Intelligence* (AAAI Press, Palo Alto, CA), 1405–1411.
- Eisenberg E, Gale D (1959) Consensus of subjective probabilities: The pari-mutuel method. *Ann. Math. Statist.* 30(1):165–168.
- Foley D (1967) Resource allocation and the public sector. *Yale Econom. Essays* 7:45–98.
- Freeman R, Zahedi SM, Conitzer V, Lee BC (2018) Dynamic proportional sharing: A game-theoretic approach. Chaintreau A, Wierman A, Akella A, eds. *Proc. ACM Measurement Anal. Comput. Systems.*, vol. 2 (ACM, New York), 1–36.
- Friedman E, Psomas CA, Vardi S (2015) Dynamic fair division with minimal disruptions. Roughgarden T, Feldman M, Schwarz M, eds. *Proc. Sixteenth ACM Conf. Econom. Comput.* (ACM, New York), 697–713.
- Friedman E, Psomas CA, Vardi S (2017) Controlled dynamic fair division. Daskalakis D, Babaioff M, Moulin H, eds. *Proc. 2017 ACM Conf. Econom. Comput.* (ACM, New York), 461–478.
- Gal Y, Mash M, Procaccia AD, Zick Y (2017) Which is the fairest (rent division) of them all? Sierra C, ed. *Proc. 26th Internat. Joint Conf. Artificial Intelligence (IJCAI, Marina del Rey, CA)*, 4841–4843.
- Gao Y, Peysakhovich A, Kroer C (2021) Online market equilibrium with application to fair division. *Adv. Neural Inform. Processing Systems* 34:27305–27318.
- Garg J, Végh LA (2019) A strongly polynomial algorithm for linear exchange markets. Charikar M, Cohen E, eds. *Proc. 51st Annual ACM SIGACT Sympos. Theory Comput.* (ACM, New York), 54–65.
- Garg J, Hoefer M, Mehlhorn K (2018) Approximating the Nash social welfare with budget-additive valuations. Czumaj A, ed. *Proc. Twenty-Ninth Annual ACM-SIAM Sympos. Discrete Algorithms* (Society for Industrial and Applied Mathematics, Philadelphia, PA), 2326–2340.
- Ghods A, Zaharia M, Hindman B, Konwinski A, Shenker S, Stoica I (2011) Dominant resource fairness: Fair allocation of multiple resource types. Andersen DG, Ratnasamy S, eds. *Proc. 8th USENIX Conf. Networked Systems Design Implementation (NSDI)* (USENIX Association, Berkeley, CA), 24–37.
- Gkatzelis V, Psomas A, Tan X (2021) Fair and efficient online allocations with normalized valuations. Yang Q, Leyton-Brown K, Mausam, eds. *Proc. 35th AAAI Conf. Artificial Intelligence (AAAI) No. 6* (AAAI Press, Palo Alto, CA), 5440–5447.
- He J, Procaccia AD, Psomas C, Zeng D (2019) Achieving a fairer future by changing the past. Kraus S, ed. *Proc. 28th Internat. Joint Conf. Artificial Intelligence (IJCAI)* (IJCAI, Marina del Rey, CA), 343–349.
- Hoeffding W (1963) Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.* 58(301):13–30.
- Kash I, Procaccia AD, Shah N (2014) No agent left behind: Dynamic fair division of multiple resources. *J. Artificial Intelligence Res.* 51:579–603.
- Kurokawa D, Procaccia AD, Wang J (2016) When can the maximin share guarantee be guaranteed? Schuurmans D, ed. *Proc. 30th AAAI Conf. Artificial Intelligence (AAAI)*, (AAAI Press, Palo Alto, CA), 523–529.
- Lee MK, Kusbit D, Kahng A, Kim JT, Yuan X, Chan A, See D, et al. (2019) WeBuildAI: Participatory framework for algorithmic governance. Lampinen A, Gergely D, Shamma DA, eds. *Proc. ACM Human Comput. Interaction* 3:1–35.
- Li B, Li W, Li Y (2018) Dynamic fair division problem with general valuations. *Proc. 27th Internat. Joint Conf. Artificial Intelligence* (ACM, New York), 375–381.
- Lipton RJ, Markakis E, Mossel E, Saberi A (2004) On approximately fair allocations of indivisible goods. Breese J, Feigenbaum J, Seltzer M, eds. *Proc. 6th ACM Conf. Econom. Comput. (EC)* (ACM, New York), 125–131.
- Mehta A, Saberi A, Vazirani U, Vazirani V (2007) Adwords and generalized online matching. *J. ACM* 54(5):22.
- Miller HE, Pierskalla WP, Rath GJ (1976) Nurse scheduling using mathematical programming. *Oper. Res.* 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. *Proc. Algorithmic Game Theory 14th Internat. Sympos. SAGT 2021* (Springer, Cham, Switzerland), 376–390.
- Orlin JB (2010) Improved algorithms for computing Fisher’s market clearing prices. Mitzenmacher M, Schulman LJ, eds. *Proc. Forty-Second ACM Sympos. Theory Comput.* (ACM, New York), 291–300.
- Procaccia AD (2016) Cake cutting algorithms. Brandt F, Conitzer V, Endress U, Lang J, Procaccia AD, eds. *Handbook of Computational Social Choice* (Cambridge University Press, Cambridge, UK), 311–330.

- Ruhe G, Fruhwirth B (1990) ε -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. Rajasekaran S, ed. *Proc. 1st Workshop Randomized Parallel Algorithms (RANDOM)*.
- Su FE (1999) Rental harmony: Sperner's lemma in fair division. *Amer. Math. 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 Sci.* 52(11):1647–1660.
- Vardi S, Psomas A, Friedman E (2022) Dynamic fair resource division. *Math. Oper. Res.* 47(2):945–968.
- Varian H (1974) Equity, envy and efficiency. *J. Econom. Theory* 9(1):63–91.
- Walsh T (2011) Online cake cutting. Perny P, Pirlot M, Tsoukifis A, eds. *Proc. 3rd Internat. Conf. Algorithmic Decision Theory (ADT)* (Springer-Verlag, Berlin, Heidelberg), 292–305.

Gerdus Benadè is an assistant professor of information systems at Boston University. He is interested in computational social choice,

fair division, and discrete optimization, especially real-world applications, like participatory budgeting, gerrymandering, and food rescue.

Aleksandr M. Kazachkov is an assistant professor in the Department of Industrial & Systems Engineering at the University of Florida. His research is on the methodology and applications of discrete optimization: in particular, integer programming techniques and designing fair mechanisms for resource allocation.

Ariel D. Procaccia is the Gordon McKay Professor of Computer Science at Harvard University. He works on a broad and dynamic set of problems related to artificial intelligence, algorithms, economics, and society, including food rescue, citizens' assemblies, and refugee resettlement.

Alexandros Psomas is an assistant professor in the Department of Computer Science at Purdue University. He is broadly interested in algorithmic economics, artificial intelligence, computational social choice and mechanism design, machine learning, and algorithmic game theory.

David Zeng works on building high-frequency crypto trading strategies. Before that, he graduated from Carnegie Mellon University, where he worked on algorithmic game theory.