



HAL
open science

Breaking the Envy Cycle: Best-of-Both-Worlds Guarantees for Subadditive Valuations

Michal Feldman, Simon Mauras, Vishnu V Narayan, Tomasz Ponitka

► **To cite this version:**

Michal Feldman, Simon Mauras, Vishnu V Narayan, Tomasz Ponitka. Breaking the Envy Cycle: Best-of-Both-Worlds Guarantees for Subadditive Valuations. EC 2024 - 25th ACM Conference on Economics and Computation, Jul 2024, New Haven, United States. pp.1236 - 1266, 10.1145/3670865.3673592 . hal-04923223

HAL Id: hal-04923223

<https://hal.science/hal-04923223v1>

Submitted on 31 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright



Breaking the Envy Cycle: Best-of-Both-Worlds Guarantees for Subadditive Valuations

MICHAL FELDMAN*, Tel Aviv University, Israel

SIMON MAURAS†, INRIA, FairPlay joint team, France

VISHNU V. NARAYAN, Tel Aviv University, Israel

TOMASZ PONITKA, Tel Aviv University, Israel

We study best-of-both-worlds guarantees for the fair division of indivisible items among agents with subadditive valuations. Our main result establishes the existence of a random allocation that is simultaneously ex-ante $\frac{1}{2}$ -envy-free, ex-post $\frac{1}{2}$ -EFX and ex-post EF1, for every instance with subadditive valuations. We achieve this result by a novel polynomial-time algorithm that randomizes the well-established envy cycles procedure in a way that provides ex-ante fairness. Notably, this is the first best-of-both-worlds fairness guarantee for subadditive valuations, even when considering only EF1 without EFX.

CCS Concepts: • **Theory of computation** → *Algorithmic game theory and mechanism design*.

Additional Key Words and Phrases: Fair Division, Envy-Cycles, Best-of-Both-Worlds, Subadditive

ACM Reference Format:

Michal Feldman, Simon Mauras, Vishnu V. Narayan, and Tomasz Ponitka. 2024. Breaking the Envy Cycle: Best-of-Both-Worlds Guarantees for Subadditive Valuations. In *Conference on Economics and Computation (EC '24), July 8–11, 2024, New Haven, CT, USA*. ACM, New York, NY, USA, 31 pages. <https://doi.org/10.1145/3670865.3673592>

1 Introduction

In a fair division problem, the goal is to allocate a set M of m indivisible items among n agents *fairly*. Every agent has a valuation function $v_i : 2^M \rightarrow \mathbb{R}^{\geq 0}$, mapping every bundle of items $S \subseteq M$ to a (non-negative) real number $v_i(S)$. An allocation $X = (X_1, \dots, X_n)$ is a partition of the items among the agents, where X_i is the bundle allocated to agent i .

The problem of allocating resources fairly dates back to Aristotle [Chroust, 1942]. Many mathematical notions of fairness have been considered in the literature (both for divisible and indivisible items). A particularly natural fairness notion that generated significant interest in the literature is *envy-freeness* (EF), requiring that every agent prefers her own bundle to any other agent's bundle [Foley, 1967, Varian, 1974]. That is, an allocation X is EF if for every two agents i and j , it holds that $v_i(X_i) \geq v_i(X_j)$. The notion of envy-freeness extends to *random* allocations in two ways. We say that a random allocation is EF *ex-ante* (before the randomization is realized) if every agent prefers

* Also with Microsoft ILDC.

† Also with Tel Aviv University.

Authors' Contact Information: Michal Feldman, mfeldman@tauex.tau.ac.il, Tel Aviv University, Israel; Simon Mauras, simon.mauras@inria.fr, INRIA, FairPlay joint team, France; Vishnu V. Narayan, narayanv@tauex.tau.ac.il, Tel Aviv University, Israel; Tomasz Ponitka, tomaszp@mail.tau.ac.il, Tel Aviv University, Israel.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

EC '24, July 8–11, 2024, New Haven, CT, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0704-9/24/07

<https://doi.org/10.1145/3670865.3673592>

her own bundle to any other agent's bundle *in expectation*. We say that a random allocation is EF *ex-post* (after the randomization is realized) if every deterministic allocation in its support is EF.

Achieving ex-ante EF is quite easy. For example, allocating all items to a single agent chosen uniformly at random is trivially ex-ante EF. However, an ex-ante EF allocation may still be arbitrarily unfair ex-post. Indeed, in the example above, one agent receives all the items, and will surely be envied by all other agents. On the other hand, ex-post EF is too strong of a requirement, as even the simplest setting of a single item desired by two agents does not admit any EF allocation.

Best-of-both-worlds fairness. Recently, Aziz [2019] posed, as an interesting new research direction, the question of finding random allocations that simultaneously achieve desirable ex-ante and ex-post properties. Shortly thereafter, Freeman et al. [2020] studied this problem in the fair division domain, with the goal of obtaining a random allocation that is simultaneously ex-ante EF and ex-post “relaxed EF”. The approach of constructing random allocations with strong ex-ante and ex-post fairness guarantees has since been known as the “best-of-both-worlds” approach.

In their work, Freeman et al. [2020] focus on the notion of EF1 – envy-freeness up to one item [Budish, 2011, Lipton et al., 2004] – which requires that the envy of every agent i toward another agent j can be removed by the elimination of at most one item from agent j 's bundle, i.e., there exists an item $g \in X_j$ such that $v_i(X_i) \geq v_i(X_j \setminus \{g\})$. Unlike EF, an EF1 allocation always exists [Budish, 2011]. Moreover, Freeman et al. [2020] and Aziz [2020] showed that any instance with additive valuations – where an agent's value for a bundle is the sum of her values for the individual items – admits a random allocation that is simultaneously ex-ante EF and ex-post EF1. Namely, an envy-free distribution over deterministic allocations such that each satisfies EF1.

While this is a promising result, it is limited in two ways. First, it is restricted to additive valuations, while in most practical settings, agent valuations are non-additive. For example, items may exhibit substitutability, where an agent might like to get one of two given items but have no additional value for having both of them. Second, this result holds with respect to EF1, but does not hold with respect to another important (and stronger) relaxation of envy-freeness, known as EFX – envy-freeness up to *any* item [Caragiannis et al., 2016] – which requires that the envy of every agent i toward another agent j can be removed by the elimination of *any* item from agent j 's bundle.

To demonstrate why EFX may be more desirable than EF1, consider the following scenario. Suppose there are three items, a, b, c , and two agents with identical additive values of 100, 50, 50 for items a, b, c , respectively. Any reasonable fairness notion in this case would allocate item a to one agent and items b, c to the other agent, resulting in a value of 100 to each agent. This allocation is the only EFX allocation for this instance. However, one can easily verify that allocating items a, b to agent 1 (for a value of 150) and item c to agent 2 (for a value of 50) is EF1: removing item a from agent 1's bundle removes agent 2's envy. Indeed, extending best-of-both-worlds results to ex-post EFX appeared as one of the open problems in Freeman et al. [2020].

While many attempts have been made towards improving the results of Freeman et al. [2020] and Aziz [2020], none of the subsequent works specifically resolve these issues. Instead, they either give best-of-both-worlds results that hold only for highly structured valuation classes, such as binary additive valuations (Aleksandrov et al. [2015], Halpern et al. [2020]), matroid rank valuations (Babaioff et al. [2021]), and multi-demand valuations (Hoefler et al. [2023]), or study the case of arbitrary entitlements, where they obtain only weaker ex-post guarantees (Aziz et al. [2023a], Hoefler et al. [2023]).

Our contribution. With the desire to go beyond additive valuations and beyond the EF1 notion, the key question we address in this paper is whether there always exists a random allocation that is simultaneously ex-ante EF and ex-post EFX, for any instance with subadditive valuations.

A valuation function is subadditive if $v(S) + v(T) \geq v(S \cup T)$ for every two bundles of items S and T . The class of subadditive valuations encompasses important subclasses, such as additive valuations and submodular valuations (namely, valuations exhibiting decreasing marginal values).

Let us consider first the simplest scenario of two agents with additive valuations. While not immediate, we show that this scenario always admits a random allocation that is simultaneously ex-ante EF and ex-post EFX (see Proposition A.3).

But this positive result breaks as soon as we go either beyond two agents or beyond additive valuations. Going beyond two agents, even for additive valuations, the mere existence of ex-post EFX, without any additional requirement, is an intriguing open problem. In fact, it is widely considered to be the biggest open problem in fair division [Procaccia, 2020]¹. Similarly, going beyond additive valuations, we show that even with only two agents and submodular valuations, there are instances that do not admit any allocation that is simultaneously ex-ante EF and ex-post EFX. This is demonstrated in Proposition A.8.

As is standard in the literature, to address this problem we turn to the framework of *approximate fairness*. A deterministic allocation is β -EFX if every agent prefers her own bundle to a β fraction of any other agent's bundle after eliminating any single item from it. A random allocation is ex-post β -EFX if every allocation in its support is β -EFX. Similarly, a random allocation is ex-ante α -EF if every agent prefers her own bundle to an α fraction of every other agent's bundle, in expectation.

The approach taken by Freeman et al. [2020] and Aziz [2020] gives no ex-post β -EFX guarantees for any $\beta > 0$ (see Appendix C). Therefore, the question that drives us in this work is the following:

Main Question: Are there constants α and β such that an allocation that is simultaneously ex-ante α -EF and ex-post β -EFX is guaranteed to exist for every profile of subadditive valuations?

Our main result is an affirmative answer to this question:

Main Result (Theorem 3.1): For every instance with subadditive valuations, Algorithm 1 outputs a random allocation that is ex-ante $\frac{1}{2}$ -EF, ex-post $\frac{1}{2}$ -EFX, and ex-post EF1 in polynomial time.

This result constitutes a significant advancement in the best-of-both-worlds fairness literature in multiple ways:

- (1) We are the first to provide best-of-both-worlds results for *subadditive* valuations – our work even extends the previous work on (EF, EF1) from additive all the way to subadditive, and gains ex-post $\frac{1}{2}$ -EFX in the process, while losing only a factor of 2 in the ex-ante envy.
- (2) No better approximation of EFX than $\frac{1}{2}$ -EFX is known to exist [Plaut and Roughgarden, 2018], even for the smaller class of submodular valuations, and even without any ex-ante EF requirements. We match this guarantee for the broader class of subadditive valuations while providing ex-ante guarantees in addition.
- (3) To the best of our knowledge, we are the first to consider randomized envy cycles in a way that provides ex-ante EF guarantees. Envy cycles is one of the most widely-used procedures for discrete fair division. Indeed, the envy cycles algorithm was one of the earliest works in the fair division of indivisible items [Lipton et al., 2004], and it has since been used to

¹EFX is only known to exist for some special cases, including three additive agents [Chaudhury et al., 2020a] or identical valuations [Plaut and Roughgarden, 2018].

achieve many fairness results, including EF1 and EFX guarantees [Lipton et al., 2004, Plaut and Roughgarden, 2018]; see Section 1.2 for more details.

We complement our positive result with the following upper bounds (holding even for 2 agents):

Impossibility Result (Proposition A.7): For every $0.618 \approx \varphi - 1 < \beta \leq 1$, there exists an instance with two subadditive valuations that admits no randomized allocation that is simultaneously ex-ante α -EF and ex-post β -EFX, for any $\alpha > \frac{\beta+1}{\beta^2+2\beta}$.

In particular, our impossibility result implies the following upper bounds:

- There is no random allocation that is ex-ante α -EF and ex-post EFX for $\alpha > \frac{2}{3}$.
- There is no random allocation that is ex-ante EF and ex-post β -EFX for $\beta > \varphi - 1$.

In Appendix A we show that the first of these bounds is tight for two agents, namely, we devise an algorithm that gives an ex-ante $\frac{2}{3}$ -EF and ex-post EFX allocation for every instance with two subadditive valuations (Proposition A.6). In addition, for every such instance, we devise an algorithm that gives an ex-ante EF and ex-post $\frac{1}{2}$ -EFX allocation (Proposition A.5).

Our results are summarized in Figure 1.

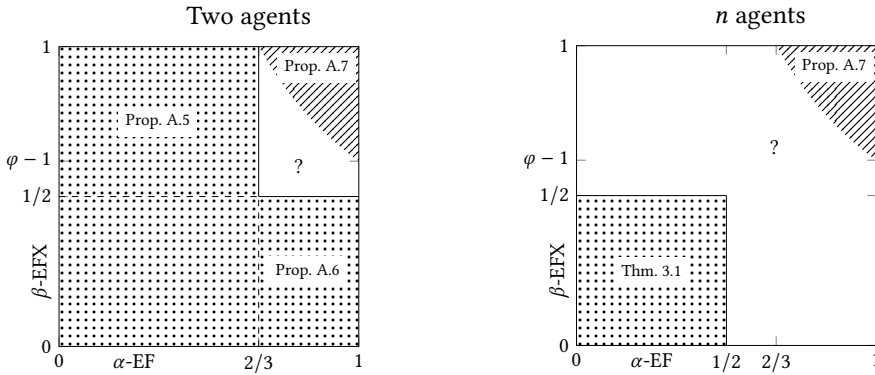


Fig. 1. Trade-offs between ex-ante α -EF and ex-post β -EFX, for two (left) and n (right) subadditive agents. The x -axis (resp., y -axis) represents ex-ante α -EF (resp., ex-post β -EFX) guarantees. Dotted (resp., hatched) areas represent our existence (resp., impossibility) results.

Discussion. While our main theorem is the first best-of-both-worlds fairness result for subadditive valuations, there are several important questions that this work raises. In Section 4, we discuss a collection of future directions for research in this area.

1.1 Our Techniques

In this section, we present an overview of the construction and the analysis of our main algorithm. While most previous works on best-of-both-worlds fairness (e.g. [Aziz, 2020, Aziz et al., 2023a, Freeman et al., 2020, Hoefer et al., 2023]) first select a suitable fractional allocation and then construct a distribution over integral allocations that implements it, our new method approaches this problem from the opposite direction. Namely, we first select a deterministic procedure with the desired ex-post properties (in this case, a simple variation of the envy cycles procedure [Lipton et al., 2004]), and then introduce a technique to randomize this procedure to obtain the ex-ante guarantees.

The standard deterministic envy cycles procedure involves two phases. In the first phase, the algorithm assigns one item to each agent using a matching algorithm that satisfies a natural property called weak separation (Definition 3.2), which states that every agent prefers her item to any unassigned item. In the second phase, the algorithm allocates the remaining items. This is done by either (1) allocating one of the remaining items to an *unenvied* agent, or (2) shifting the bundles along an *envy cycle*, which is a cycle where each agent prefers the next agent’s bundle to her own. This procedure is explained in detail in Section 3.1.

Crucially, there are many ways to implement the deterministic envy cycles procedure that yield the same ex-post guarantees. For instance, in the second phase, one can arbitrarily decide which unenvied agent to select for operation (1) or which envy cycle to choose for operation (2). However, for *any* deterministic choices for these operations, this algorithm gives no ex-ante EF guarantees. Consequently, achieving the desired ex-ante fairness requires randomization of these choices. In what follows, we describe the randomization of each phase of the algorithm separately.

Randomizing the First Phase. One possible approach to obtain ex-ante fairness guarantees is to apply some randomness to a given deterministic algorithm. For example, a common algorithm for the first phase is serial dictatorship, where agents are sorted according to some order, and every agent, upon her turn, chooses her most preferred item among the remaining ones. One can easily verify that deterministic serial dictatorship gives no ex-ante fairness guarantees. A natural approach to obtain some ex-ante fairness guarantees is to randomize the order in which agents choose items. This is called random serial dictatorship (RSD). Unfortunately, choosing a uniformly random order over the agents does not give ex-ante EF [Bogomolnaia and Moulin, 2001]. In fact, we improve the upper bound on the ex-ante fairness of RSD, by showing (Proposition D.2) that using a uniformly random order does not give better than ex-ante $\frac{1}{\sqrt{2}}$ -envy-freeness.

Therefore, to construct the matching in the first phase, we use the *probabilistic serial* procedure of Bogomolnaia and Moulin [2001]. Recently, Freeman et al. [2020] and Aziz [2020] utilized this procedure to achieve ex-ante EF together with ex-post EF1 for additive valuations. They apply the probabilistic serial procedure for the entire allocation process. Unfortunately, this method does not extend to subadditive valuations, and even for additive valuations, it does not guarantee β -EFX for any $\beta > 0$ (see Appendix C). Instead, we apply the probabilistic serial method for a single time unit, followed by randomized envy cycles. See Section 3.2 for more details.

Randomizing the Second Phase. It may seem reasonable to hope that selecting the envy cycles arbitrarily during the second phase of the algorithm results in a fair outcome. However, even though an envy cycle elimination step does not modify any of the bundles, it might shuffle them among the agents in a way that benefits some agents more than others. This is illustrated in Example 3.11. A natural second attempt would be to choose the envy cycle uniformly at random. However, one can show that this too does not provide any ex-ante EF guarantees.

We propose a novel way to sample the envy cycles in a way that provides a probability distribution that gives ex-ante EF guarantees. To construct the desired distribution, we present an algorithm inspired by the decomposition of irreducible Markov chains into circuit processes by MacQueen [1981]. This algorithm applies only to strongly connected graphs, and to address this challenge, we first select a strongly connected component of the envy graph that has no incoming edges (which always exists), and run the algorithm on the induced subgraph. Details regarding the randomization of the second phase are given in Section 3.3.

Following the explanation of the main algorithm’s design, we now focus on the components of the analysis of its ex-ante fairness guarantees. We first identify two key properties of the algorithm.

First, the random allocation obtained in the first phase satisfies a property that we term *strong separation* (see Lemma 3.10). This property states that every agent prefers her bundle in any allocation in the support of the random allocation to any item that may be unallocated in any (possibly different) allocation in the support. Second, we show that after any step of the algorithm, the value of every agent i for her own bundle *stochastically dominates* the value of i for (a specific part of) any other agent's bundle (see Lemma 3.15).

To establish our main result, we first use the stochastic dominance property (Lemma 3.15) to show that *two copies* of X_i can “cover” a specific subset of X_j , in a way that is captured by the definition of *stochastic coverage* that we introduce in Section 3.4. The proof is by induction on the execution tree, i.e., the tree representing all possible random choices that the algorithm makes on a given input. Moreover, using strong separation (Lemma 3.10), we extend this property to show that the two copies of X_i can also cover the remaining part of X_j , from which our main result follows easily. The full analysis of the main algorithm is given in Section 3.4.

1.2 Related Work

While the fair division problem has been central to human society since antiquity, its formal study began with the work of Banach, Knaster and Steinhaus [Steinhaus, 1948], who analyzed the *cake-cutting* problem: *how can a heterogeneous cake be divided fairly amongst agents?* With $n = 2$ agents, the folklore *cut-and-choose* method results in a *proportional* allocation of the cake, i.e., one in which each agent is allocated a piece of value at least $\frac{1}{n}$ of their value ($\frac{1}{2}$ in the case of $n = 2$) for the whole cake. This result was extended to $n > 2$ agents with the *last diminisher* procedure [Steinhaus, 1948]. Subsequent decades witnessed the emergence of *envy-freeness* (Foley [1967], Varian [1974]) as the main criterion for fairness in economics. An allocation is envy-free if each agent prefers her own bundle to the bundle of any other agent. For the settings with divisible items, which generalize the cake-cutting problem, envy-free allocations are known to exist under mild assumptions (Stromquist [1980]), and many algorithms have been devised for their computation [Aziz and Mackenzie, 2016, Brams and Taylor, 1995].

The Envy Cycles Procedure. For the complementary setting with *indivisible* items, both proportionality and envy-freeness are impossible to achieve: consider the simple instance with two agents and one item, where any allocation leaves one agent envying the other. A breakthrough development in the area occurred almost two decades ago, when Lipton et al. [2004] introduced the now-ubiquitous *envy cycles procedure* for the fair division of indivisible items. This led to the development of relaxed fairness notions such as the EF1 guarantee, formally defined by Budish [2011]. The envy cycles procedure computes, in polynomial time, an EF1 allocation even for the very general class of monotone valuations. This work subsequently spawned a large body of research in the fair division of indivisible items, and the envy cycles procedure (along with its variations) has since been used in a wide variety of studies in fair division, including in algorithms for finding approximate MMS allocations [Barman and Krishnamurthy, 2020], EF1 allocations for chore division [Bhaskar et al., 2021], approximate EFX allocations [Plaut and Roughgarden, 2018], and partial EFX allocations (“EFX with charity”) [Chaudhury et al., 2020b], among several others.

The EFX and MMS Fairness Notions. The EFX guarantee was introduced by Caragiannis et al. [2016]. Despite its apparent similarity to EF1, the existence of EFX allocations remains a notoriously hard open problem even for additive agents. Complete EFX allocations are known to exist for instances with at most three additive agents (Chaudhury et al. [2020a]), and in the case where the agents have identical valuation functions (Plaut and Roughgarden [2018]). Beyond these results, only improvements for special cases or for various relaxations have been found [Akrami et al., 2022,

Ghosal et al., 2023, Mahara, 2021, Plaut and Roughgarden, 2018]. Because of the limited progress on the existence of complete EFX allocations, both *partial*- and *approximate*-EFX allocations have also been studied. Numerous works have studied the existence of partial allocations that provide some EFX guarantee while allocating almost all of the items [Akrami et al., 2022, Berger et al., 2022, Chaudhury et al., 2021, 2020b, Feldman et al., 2024]. For approximately-EFX fair division, Plaut and Roughgarden [2018] showed that a $\frac{1}{2}$ -EFX allocation exists even for subadditive valuations. For additive valuations, Amanatidis et al. [2020] improved this approximation factor to $\varphi - 1 \approx 0.618$.

Another related fairness notion for the indivisible setting, defined by Budish [2011], is the maximin share (MMS) guarantee. It is known that allocations that give each agent its MMS value do not exist even for additive agents (Feige et al. [2021], Procaccia and Wang [2014]), but there are allocations that give each agent a constant factor of the MMS value, and these allocations can be efficiently computed even for submodular and XOS valuations (see, e.g., [Barman and Krishnamurthy, 2020, Ghodsi et al., 2018]).

Best-of-Both-Worlds Fairness. A recent line of research, titled best-of-both-worlds fairness, aims to achieve envy-freeness for indivisible items via randomization. The leading question is whether it is possible to simultaneously achieve some ex-ante fairness guarantee while ensuring ex-post guarantees for every realized outcome of the random process. Aleksandrov et al. [2015] first studied this problem in a food bank setting, where they gave an algorithm with ex-ante and ex-post fairness guarantees for the special case of additive valuations with *binary (0/1)* marginals. For additive valuations, Freeman et al. [2020] gave a randomized polynomial-time algorithm that outputs an EF1 allocation, while being envy-free ex-ante. Ensuing work by Aziz [2020] showed that there exists a similar randomized algorithm that additionally implements the well-known *Probabilistic Serial* fractional outcome described by Bogomolnaia and Moulin [2001]. This randomized algorithm also preserves a weak notion of efficiency. Babaioff et al. [2022] also study the case of additive valuations and find a distribution over ex-post proportional up to one item and $\frac{1}{2}$ -MMS allocations that is ex-ante proportional, i.e., the expected value of each agent's bundle is at least a $\frac{1}{n}$ -fraction of her value for the set of all items.

Best-of-both-worlds results have also been analyzed for other settings, including the case where the agents have binary marginals, and the case of additive valuations with *arbitrary entitlements*. For the special case of additive valuations with binary marginals, the algorithm of Aziz [2020] improves upon the guarantees of Aleksandrov et al. [2015] and is group-strategyproof, ex-ante fractionally-PO and envy-free, and ex-post fractionally PO and EF1. In a similar vein, for additive valuations with binary marginals, Halpern et al. [2020] independently showed that there is a distribution over ex-post Nash-welfare-maximizing allocations that also ex-ante maximizes the fractional Nash welfare, implying the same fairness guarantees as Aziz [2020] for this setting. For *matroid rank* valuations, Babaioff et al. [2021] present a randomized truthful mechanism that is ex-ante envy-free and ex-post Lorenz dominating (and thus ex-post Nash-welfare-maximizing and EFX for this specific class). For the case of agents with arbitrary entitlements, both Hoefer et al. [2023] and Aziz et al. [2023a] show that weighted ex-ante envy-freeness along with ex-post weighted envy-freeness up to 1 item is impossible to achieve. Both Hoefer et al. [2023] and Aziz et al. [2023a] then give polynomial time algorithms that instead achieve some weaker fairness properties. Additionally, Hoefer et al. [2023] show that ex-ante EF and ex-post EF1 can be achieved for multi-demand valuations. Finally, Akrami et al. [2023] study best-of-both-worlds fairness for the MMS guarantee for XOS valuations.

Beyond item allocation, best-of-both-worlds approaches have been utilized in other areas, such as committee voting [Aziz et al., 2023b] and participatory budgeting [Aziz et al., 2024].

2 Preliminaries

An instance of the resource allocation problem consists of a set $[n] = \{1, \dots, n\}$ of n agents, a set $[m] = \{1, \dots, m\}$ of m indivisible items, and a valuation profile $(v_i)_{i \in [n]}$. The valuation function $v_i : 2^{[m]} \rightarrow \mathbb{R}^{\geq 0}$ of agent i gives a non-negative real value $v_i(S)$ for every bundle of items $S \subseteq [m]$. We assume that all valuation functions are monotone, i.e., v_i satisfies $v_i(S) \leq v_i(T)$ for any $S \subseteq T \subseteq [m]$ and every agent i . The valuation functions we consider in this work are either

- additive, i.e., it holds that $v_i(S) = \sum_{x \in S} v_{ix}$ for some $v_{i1}, \dots, v_{im} \geq 0$, or
- subadditive, i.e., it holds that $v_i(S \cup T) \leq v_i(S) + v_i(T)$ for any two $S, T \subseteq [m]$.

We use the following standard notation for singleton sets: for any $g \in [m]$ and $S \subseteq [m]$, we write $v_i(g) = v_i(\{g\})$ and $S + g = S \cup \{g\}$ and $S - g = S \setminus \{g\}$.

A *deterministic* partial allocation $X = (X_1, \dots, X_n)$ of items to agents is a partition of the items into n bundles where all bundles are disjoint, i.e., it holds that $X_i \cap X_j = \emptyset$ for all $i \neq j$. We say that X is *complete* if no items are unallocated, i.e., it holds that $\bigcup_{i \in [n]} X_i = [m]$.

We say that a deterministic allocation X is α -EF1 (for some $\alpha \in [0, 1]$) if for all agents i and j , it holds that $v_i(X_i) \geq \alpha \cdot v_i(X_j - g)$ for *some* item $g \in X_j$. We say that an allocation X is α -EFX (for some $\alpha \in [0, 1]$) if it holds that $v_i(X_i) \geq \alpha \cdot v_i(X_j - g)$ for *all* items $g \in X_j$. An allocation is EF1 (resp., EFX) if it is 1-EF1 (resp., 1-EFX).

Random allocations, ex-post and ex-ante fairness. A *random* allocation $X = (X_1, \dots, X_n)$ is a random variable X whose every random outcome is a deterministic allocation. Here, X_1, \dots, X_n are the associated random variables whose every random outcome is a bundle of items. A random allocation X is *ex-post* α -EF1 (resp., α -EFX) if every random outcome of X (that happens with non-zero probability) is an α -EF1 (resp., α -EFX) deterministic allocation. We say that X is *ex-ante* α -EF (for some $\alpha \in [0, 1]$) if $\mathbb{E}[v_i(X_i)] \geq \alpha \cdot \mathbb{E}[v_i(X_j)]$ for all i, j . An allocation is EF if it is 1-EF.

Envy graph and strongly connected components. For every deterministic partial allocation $X = (X_1, \dots, X_n)$, we define the envy graph $G_X = ([n], E_X)$ where the set of nodes is the set of agents $[n]$ and the set of directed edges $E_X = \{(i, j) \in [n]^2 : v_i(X_j) > v_i(X_i)\}$ includes an edge from agent i to agent j if i envies j , i.e., agent i prefers agent j 's bundle over her own bundle.

A directed graph is *strongly connected* if there exists a directed path from any node to any other node. A *strongly connected component* of a directed graph G is any maximal (in the sense of inclusion) subgraph of G that is strongly connected.

Stochastic dominance. Let X and Y be any two non-negative real random variables. We say that X *stochastically dominates* Y , and we write $X \succeq_{SD} Y$, if it holds that $\mathbb{P}[X \geq t] \geq \mathbb{P}[Y \geq t]$ for every $t \geq 0$.

3 Our Main Result: Ex-Ante $\frac{1}{2}$ -EF, Ex-Post $\frac{1}{2}$ -EFX, Ex-Post EF1

In this section, we describe our main result: the existence of a distribution over $\frac{1}{2}$ -EFX and EF1 allocations that is ex-ante $\frac{1}{2}$ -EF for every instance with subadditive valuations, and a randomized algorithm that outputs a random allocation with this distribution in polynomial time. This result is presented in Theorem 3.1, which we restate below.

THEOREM 3.1. *For every instance with subadditive valuations, Algorithm 1 outputs a random allocation that is ex-ante $\frac{1}{2}$ -EF, ex-post $\frac{1}{2}$ -EFX, and ex-post EF1 in polynomial time.*

We may assume without loss of generality that there are more than n items. Indeed, if there are $m \leq n$ items, then we may add $n - m$ dummy items with value 0 for all agents, and choose

a random allocation that assigns one item per agent. The resulting allocation is EFX (since each agent is assigned at most one item) and ex-ante EF (by symmetry).

3.1 A Deterministic Envy Cycles Procedure

In this section we describe a deterministic algorithm for constructing $\frac{1}{2}$ -EFX and EF1 allocations (but, with no EF guarantees) when the agents have subadditive valuations. As we will show in the following subsections, our main algorithm is a carefully randomized implementation of this deterministic algorithm.

The existence of $\frac{1}{2}$ -EFX and EF1 allocations for subadditive agents was first shown by Plaut and Roughgarden [2018]. Their proof is constructive and involves the use of an algorithm that is based on the envy cycles procedure. We have modified the presentation of the algorithm to suit our analysis, and we present it here as a deterministic two-phase algorithm that runs in polynomial time, see Algorithm 2 in Appendix B.

In the first phase, the algorithm constructs a *matching* between the agents and the items, in which exactly one item is assigned to each agent. The matching must satisfy the following property, called *weak separation*: each agent prefers the item she receives in the matching to any of the remaining unassigned items. The formal definition of this property was explicitly given in Feldman et al. [2024], but was implicitly used in previous studies on EFX allocations [Amanatidis et al., 2020, Barman et al., 2018, Chan et al., 2019].

Definition 3.2 (Weak separation). A partial allocation $X = (X_1, \dots, X_n)$ with the set $U = [m] \setminus \bigcup_{i \in [n]} X_i$ of unallocated items satisfies *weak separation* if for every agent $i \in [n]$ and every unallocated item $u \in U$ it holds that $v_i(X_i) \geq v_i(u)$.

In the second phase, the algorithm performs the widely-known envy cycles procedure of Lipton et al. [2004] to allocate the remaining unassigned items, as follows. Suppose that the algorithm has already constructed a partial allocation $X = (X_1, \dots, X_n)$. Then, in the next step, if there exists an *unenvied* agent j (i.e., an agent j such that $v_i(X_i) \geq v_i(X_j)$ for all $i \neq j$), then the algorithm selects an arbitrary unallocated item g (i.e., g is not in the set $\bigcup_{i \in [n]} X_i$), and assigns item g to agent j . If there is no unenvied agent, the algorithm finds a directed cycle $C = (u_1, u_2, \dots, u_k)$ in the envy graph, where agent u_1 envies agent u_2 , agent u_2 envies agent u_3 , ..., and agent u_k envies agent u_1 . Observe that such a cycle always exists; indeed, if there is no unenvied agent, every node in the envy graph has an incoming edge. The algorithm then performs a *cycle elimination* step, in which the fixed bundles are redistributed along the cycle (namely, every agent u_i receives the bundle $X_{u_{i+1}}$, where $u_{k+1} = u_1$).

We first show the following monotonicity lemma, whose proof is deferred to Appendix B.

LEMMA 3.3. *For any agent i , the value of i for her own bundle weakly increases at each step of the envy cycles procedure.*

Next, we show that the final allocation of this algorithm is $\frac{1}{2}$ -EFX and EF1. This is shown by Plaut and Roughgarden [2018] for their algorithm, but as our algorithm is presented in a slightly different way, we provide the proof in Appendix B to ensure completeness.

LEMMA 3.4 ([PLAUT AND ROUGHGARDEN, 2018]). *The allocation returned by Algorithm 2 is $\frac{1}{2}$ -EFX and EF1.*

We also show in Appendix B that our algorithm runs in polynomial time.

LEMMA 3.5. *Algorithm 2 terminates after polynomially many steps.*

3.2 A Distribution over Matchings

For the first phase of the algorithm, we follow the approach of Bogomolnaia and Moulin [2001], who study the allocation of n items among n agents, where every agent receives exactly one item. Their algorithm applies Probabilistic Serial, followed by Birkhoff-von Neumann rounding, described below.

In Probabilistic Serial, the allocation is constructed via the *simultaneous eating* procedure, where every agent i “consumes” her favorite item, i.e., the item g that maximizes $v_i(g)$ among all available items (with ties between multiple items of the same value broken arbitrarily), at a constant rate of one item per one unit of time. During this process, all of the agents have the same eating rate. After any item is fully consumed (possibly by multiple agents), each agent independently and instantaneously switches to her next-best item (again, with ties broken arbitrarily) and continues eating at the same constant rate.

Bogomolnaia and Moulin [2001] study the setting where n items are allocated to n agents, and every agent receives exactly one item. In our case, however, the number of items m might be greater than the number of agents n , and we apply this procedure for exactly one time unit (i.e., at the point where each agent has consumed a total fractional mass of exactly one item). This is in contrast to the generalization of Freeman et al. [2020] and Aziz [2020], who repeatedly apply probabilistic serial until all m items are allocated. We refer the reader to Appendix C for a more detailed explanation of Probabilistic Serial and the algorithms of Freeman et al. [2020] and Aziz [2020].

The output of the algorithm is a fractional allocation of items to agents (see Figure 4(a)). It is represented by a matrix $Z_{ij} \in [0, 1]$ for $1 \leq i \leq n$ and $1 \leq j \leq m$, satisfying $\sum_{j \in [m]} Z_{ij} = 1$ for every agent $i \in [n]$ and $\sum_{i \in [n]} Z_{ij} \leq 1$ for every item $j \in [m]$.

We then apply a classic theorem of Birkhoff [Birkhoff, 1946] and von Neumann [Von Neumann, 1953] (which we restate here in a form that is relevant to our algorithm) to decompose the fractional allocation returned by the eating procedure into a distribution over integral allocations.

THEOREM 3.6 (BIRKHOFF-VON-NEUMANN). *Let Z be the fractional allocation returned after one time step of Probabilistic Serial. There is a strongly-polynomial-time algorithm that computes permutation matrices X^1, \dots, X^q (where $X_{ij}^k \in \{0, 1\}$ for all i, j, k , and every row and column of each X^k has exactly one 1-entry) and probabilities $p^1, \dots, p^q \in (0, 1]$ (where $\sum_{k \in [q]} p^k = 1$) such that $Z = \sum_{k \in [q]} p^k X^k$, i.e., the distribution $((X^k, p^k))_{k \in [q]}$ over integral allocations is a weighted decomposition of the fractional allocation Z .*

The above theorem implies the following corollary.

COROLLARY 3.7. *For the fractional allocation Z returned by one time step of Probabilistic Serial, and its decomposition $((X^k, p^k))_{k \in [q]}$ obtained in Theorem 3.6, we have*

- (i) *for every agent $i \in [n]$ and $k \in [q]$, it holds that $\sum_{j \in [m]} X_{ij}^k = 1$, i.e., each agent is assigned exactly one item in every integral allocation X^k in the support of the decomposition.*
- (ii) *for every item $j \in [m]$ for which $\sum_{i \in [n]} Z_{ij} = 0$, it holds that $\sum_{i \in [n]} X_{ij}^k = 0$, i.e., the item j is unallocated in every integral allocation X^k in the support of the decomposition.*
- (iii) *for every item $j \in [m]$ for which $\sum_{i \in [n]} Z_{ij} = 1$, it holds that $\sum_{i \in [n]} X_{ij}^k = 1$, i.e., the item j is allocated in every integral allocation X^k in the support of the decomposition.*

Sampling an integral allocation $X = (X_1, \dots, X_n)$ from the distribution given by Theorem 3.6 has several advantages. First, it has strong ex-ante properties, captured by the following lemma.

LEMMA 3.8 ([BOGOMOLNAIA AND MOULIN, 2001]). *For any allocation X obtained by Theorem 3.6, it holds that $v_i(X_i) \geq_{SD} v_i(X_j)$ for any two agents i and j .*

PROOF. We show that $\mathbb{P}[v_i(X_i) \geq t] \geq \mathbb{P}[v_i(X_j) \geq t]$ for any $t \geq 0$. Let $\mathcal{A} = \{a \in [m] : v_i(a) \geq t\}$. Note that $\mathbb{P}[v_i(X_i) \geq t] = \sum_{a \in \mathcal{A}} \mathbb{P}[X_i = \{a\}] = \sum_{a \in \mathcal{A}} Z_{ia}$ and $\mathbb{P}[v_i(X_j) \geq t] = \sum_{a \in \mathcal{A}} Z_{ja}$. Suppose for the purpose of contradiction that $t_1 = \sum_{a \in \mathcal{A}} Z_{ia} < \sum_{a \in \mathcal{A}} Z_{ja} = t_2$. This means that after t_1 units of time, agent i started consuming an item of value less than t . However, between time t_1 and t_2 , agent j was still consuming some items in \mathcal{A} which means that at least one of them was still not fully consumed during that time. This contradicts the assumption that i always switches to eat the next most valuable item that is not fully consumed yet. \square

Moreover, it is easy to see that any integral allocations in the support of the distribution given by Theorem 3.6 is weakly separated (Definition 3.2). Hence, the ex-post guarantees given in Section 3.1 are always satisfied.

However, to prove the desired ex-ante guarantees given in Section 3.4, weak separation is not enough, and we use the following stronger notion.

Definition 3.9 (Strong separation). A random partial allocation $X = (X_1, \dots, X_n)$ with the corresponding random collection $U = [m] \setminus \bigcup_{i \in [n]} X_i$ of unallocated items satisfies *strong separation* if it holds that for any allocation Y in the support of X , we have $v_i(Y_i) \geq v_i(w)$ for every agent $i \in [n]$ and every item w such that $\mathbb{P}[w \in U] > 0$.

Clearly, any strongly separated allocation is also weakly separated. We now show that the outcome of the first phase is strongly separated.

LEMMA 3.10. *The (random) allocation $X = (X_1, \dots, X_n)$ is strongly separated.*

PROOF. Fix any agent i . Let $a \in [m]$ be any item that agent i ate a positive fraction of, i.e., it holds that $Z_{ia} = \mathbb{P}[X_i = \{a\}] > 0$. Let $b \in [m]$ be any item that was not fully consumed during the eating process, i.e., it holds that $1 - \sum_{k \in [n]} Z_{kb} = \mathbb{P}[b \in U] > 0$. Then, it is the case that $v_i(a) \geq v_i(b)$. Otherwise, agent i would have started eating b before eating any fraction of a . \square

3.3 A Distribution over Envy Cycles

Having described the implementation of the first phase of the algorithm, let us now discuss the second phase. Suppose that the algorithm already selected a partial allocation (X_1, \dots, X_n) and at the current step, the algorithm must either (1) assign an unallocated item to an unenvied agent if there is such an agent, or (2) eliminate an envy cycle by reallocating the bundles along the cycle if there is such a cycle.

First, it is important to note that the ex-post guarantees of the deterministic algorithm given by Lemma 3.4 do not depend on the specific choices of the algorithm. The algorithm can follow an arbitrary rule to decide (a) whether to perform operation (1) or (2) if at some point both of them can be executed, (b) which unenvied agent to give a new item to, (c) which unallocated item to assign to an unenvied agent, and (d) which envy cycle to eliminate.

However, when it comes to ex-ante guarantees, arbitrary choices regarding the above dimensions can lead to an arbitrarily high ex-ante envy. This is demonstrated in the following example.

Example 3.11. Consider an instance with n agents having additive valuations over $2n + 1$ items, as described in the table below.

	apple	banana	celery	durians $\times (n - 3)$			eggplants $\times (n + 1)$		
v_1	10	3ε	2ε	0	...	0	2ε	...	2ε
v_2	10	2ε	3ε	0	...	0	2ε	...	2ε
v_3	10	8	0	9	...	9	3	...	3
v_4	10	8	0	9	...	9	3	...	3
\vdots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots	\vdots	\ddots	\vdots
v_n	10	8	0	9	...	9	3	...	3

Observe that in this example, the apple is the most-preferred item for every agent. Suppose that in the first phase of the algorithm, the following allocation is chosen: agent i (for some $3 \leq i \leq n$) gets the apple, agent 1 gets the banana, agent 2 gets celery, and each of the remaining agents $3, \dots, i - 1, i + 1, \dots, n$ gets a durian. This is, for instance, exactly the allocation that is chosen with high probability (i.e., probability $1 - o(1)$) by both the probabilistic serial method and the serial dictatorship algorithm with a uniformly random permutation of the agents.

Furthermore, suppose that in the first step of the second phase, agent 1 (who is unenvied at the moment) gets an eggplant (which is unallocated at the moment). After this operation, there is an envy cycle: agent 1 (who has the banana and an eggplant) envies agent i (who has the apple), who in turn envies agent 1. However, if the algorithm eliminates this cycle, then agent 2 (who has the celery) envies agent 1 (who now has the apple) by an arbitrarily large factor. Moreover, for sufficiently small ε , once agent 1 receives the apple she will not be included in any envy cycles in the subsequent steps, so the apple will stay with her until the termination of the algorithm.

Example 3.11 shows that during a cycle elimination step, selecting the envy cycle in an arbitrary way does not give ex-ante α -EF for any $\alpha > 0$. To overcome this problem, we employ two ideas. First, we only eliminate envy cycles if there are no unenvied agents. In Example 3.11, this means that we do not immediately allow agent 1 to trade for the apple, because agent 2 remains unenvied.

The second (and key) idea is the following. Suppose that there are two agents i_1 and i_2 who both envy some bundle X_j . We want to avoid selecting an envy cycle containing the arc (i_1, j) where i_1 gets X_j , unless, with sufficiently high probability, we also select some other envy cycle containing the arc (i_2, j) , where i_2 gets X_j . Otherwise, agent i_2 might have high ex-ante envy towards agent i_1 . The corresponding intuition in Example 3.11 is that we only allow agent 1 to trade for the apple if agent 2 also has a sufficiently high probability of getting the apple.

Now, if the envy graph is strongly connected, then applying the following key lemma to the envy graph gives a “fair” distribution over the envy cycles. The proof of this lemma is inspired by the work of MacQueen [1981] on circuit processes. Figure 2 provides an illustration for this proof.

LEMMA 3.12. (Key Lemma) *Let $G = (V, E)$ be a strongly connected directed graph with $|V| \geq 2$. There exists a probability distribution $((c_t, p_t))_{t \in [r]}$ over the set of simple cycles c_t in G such that for all $j \in V$ and $(i_1, j), (i_2, j) \in E$ it holds that $\sum_{t: (i_1, j) \in c_t} p_t = \sum_{t: (i_2, j) \in c_t} p_t$, i.e., the total probability of all cycles containing the edge (i_1, j) is equal to total probability of all cycles containing the edge (i_2, j) . Moreover, such a distribution can be computed in polynomial time.*

PROOF. Consider a Markov chain with the set of states V and the transition probabilities

$$p(j, i) = \begin{cases} 1/\text{in-deg}_j & \text{if } (i, j) \in E \\ 0 & \text{otherwise} \end{cases}$$

The probabilities p are well-defined because $\sum_{i \in V} p(j, i) = 1$ for all $j \in V$. Since G is strongly connected, this Markov chain is irreducible and therefore has a stationary distribution $\pi = (\pi_i)_{i \in V}$. Note that the stationary distribution is unique [Norris, 1998, Section 1.7], and we can therefore compute the distribution in polynomial time with a direct method such as Gaussian elimination.

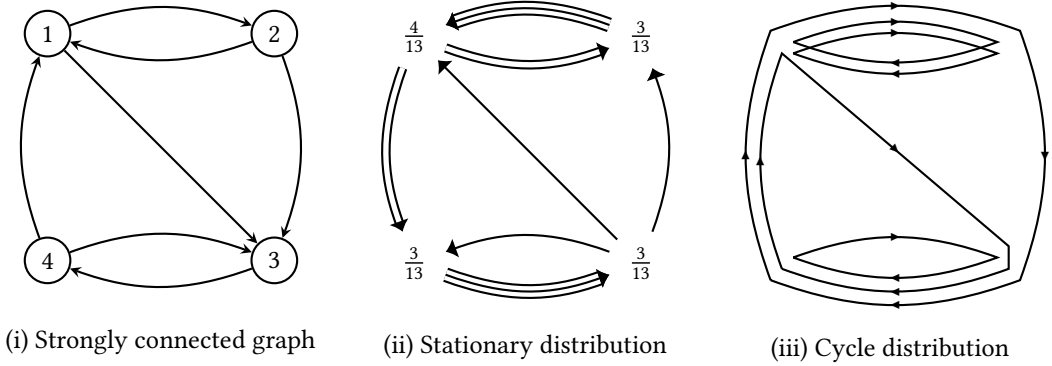


Fig. 2. Illustration of the proof of Lemma 3.12 for a strongly connected graph. The graph in (i) induces the stationary distribution in the Markov chain presented in (ii), which can be used to compute the cycle distribution presented in (iii). The multiplicity of each arrow in (ii) is proportional to its flow, as defined by w in the proof, i.e., each line represents a flow of $1/13$.

Define the flow $w(j, i) = \pi(j) \cdot p(j, i)$ and observe that w is a balanced flow, i.e., the flow that goes into i is equal to the flow that goes out of i . This is true because

$$\sum_{i \in V} w(j, i) = \pi(j) = \sum_{i \in V} \pi(i) \cdot p(i, j) = \sum_{i \in V} w(i, j),$$

where the second equality follows from the properties of the stationary distribution π .

Next, we show that there is a probability distribution $((c_t, p_t))_{t \in [k]}$ over the set of simple cycles c_t in G such that for all $(i, j) \in E$ it holds that $\sum_{t: (i,j) \in c_t} p_t = w(j, i) \cdot C$ for some constant $C > 0$. This implies that for any two edges $(i_1, j), (i_2, j) \in E$ it holds that

$$\sum_{t: (i_1, j) \in c_t} p_t \cdot C^{-1} = w(j, i_1) = \pi(j) \cdot p(j, i_1) = \pi(j) / \text{in-deg}_j = w(j, i_2) = \sum_{t: (i_2, j) \in c_t} p_t \cdot C^{-1}.$$

We construct our collection of cycles in the following manner. We start with an empty collection of cycles. In each step t , we find a cycle $((i_1, i_2), (i_2, i_3), \dots, (i_{k-1}, i_k), (i_k, i_1))$ on which every arc has positive flow (this can be done in polynomial time via, e.g., depth-first search), and then find the minimum flow w_t over any arc in this cycle. We add the pair (c_t, p_t) to our collection where $c_t = ((i_1, i_k), (i_k, i_{k-1}), \dots, (i_3, i_2), (i_2, i_1))$ and $p_t = w_t$, and then reduce the flow on the arcs $(i_1, i_2), (i_2, i_3), \dots, (i_{k-1}, i_k)$ by w_t . Note that the flow is still balanced after this operation. We repeat this process until there no longer exists a cycle with positive flow on every arc.

Since in each step of the above process, the flow on at least one arc is reduced to 0, this process terminates after a polynomial number of steps. Additionally, since w is balanced at every step, this process only terminates when the flow is zero everywhere and we thus obtain $\sum_{t: (i,j) \in c_t} p_t = w(j, i)$. Finally, we normalize the probabilities over the cycles in our collection so that they sum to 1. \square

It remains to consider the case where the envy graph is not strongly connected. In this case, we apply the key lemma to a strongly connected component G with no incoming arcs, in order to avoid generating ex-ante envy for the agents outside of G . Such a component must always exist since the connectivity graph of the strongly connected components is acyclic. Moreover, since we only eliminate envy cycles if all the agents are envied, the selected component contains at least two agents, since otherwise, the single agent in the component would be unenvied.

The distribution over the envy cycles in G given by Lemma 3.12 satisfies the following property: if any agent i prefers some bundle X_j to her own bundle X_i , then no other agent is more likely to get X_j than i . We use this property in the proof of the following lemma.

LEMMA 3.13. *Let (X'_1, \dots, X'_n) be the (random) allocation obtained after eliminating an envy cycle chosen according to the probability distribution given by Lemma 3.12 applied to the strongly connected component G . Then, for any two agents i and j , it holds that*

$$v_i(X'_i) \geq_{\text{SD}} v_i(X'_j) \cdot \mathbb{1}[X'_j \neq X_j].$$

PROOF. Let C be the (random) envy cycle chosen according to the distribution given in Lemma 3.12. First, observe that for any $t \leq v_i(X_i)$, it holds that $\mathbb{P}[v_i(X'_i) \geq t] = 1$ since $v_i(X'_i) \geq v_i(X_i) \geq t$ by Lemma 3.3, and hence $\mathbb{P}[v_i(X'_i) \geq t] \geq \mathbb{P}[v_i(X'_j) \cdot \mathbb{1}[X'_j \neq X_j] \geq t]$.

Let $t > v_i(X_i)$ and $\mathcal{K} = \{k \in [n] : v_i(X_k) \geq t\}$. Note that $v_i(X'_i) \geq t$ if and only if $X'_i = X_k$ for some $k \in \mathcal{K}$, i.e., it holds that $(i, k) \in C$. Similarly, $v_i(X'_j) \cdot \mathbb{1}[X'_j \neq X_j] \geq t$ if and only if $X'_j = X_k$ for some $k \in \mathcal{K} - j$, i.e., it holds that $(j, k) \in C$. By the properties of C given in Lemma 3.12, it holds that

$$\mathbb{P}[v_i(X'_i) \geq t] = \sum_{k \in \mathcal{K}} \mathbb{P}[(i, k) \in C] \geq \sum_{k \in \mathcal{K} - j} \mathbb{P}[(j, k) \in C] = \mathbb{P}[v_i(X'_j) \cdot \mathbb{1}[X'_j \neq X_j] \geq t]$$

which gives the result. \square

3.4 The Main Algorithm

We now present the main algorithm (Algorithm 1) that combines the ideas discussed above.

To analyze the execution of the algorithm, we first define the execution tree of all the possible random choices that the algorithm makes on a given input.

Definition 3.14 (Execution tree). Let \mathcal{H}^t be the collection of all the possible choices of the algorithm up to the t -th round, and let $\mathcal{H} = \bigcup_{t \geq 0} \mathcal{H}^t$. The set of vertices in the execution tree is \mathcal{H} and any vertex $u \in \mathcal{H}^t$ is connected with an edge to all the vertices $u' \in \mathcal{H}^{t+1}$ such that the first t rounds in u' are identical to u .

In particular, note that there is a single element in \mathcal{H}^0 , and q elements in \mathcal{H}^1 , one for each allocation in the support of the random outcome of the first phase. Every node $u \in \mathcal{H}^t$ in the execution tree is naturally associated with the probability $\mathbb{P}[u]$ that the algorithm reaches the state u , i.e., the realization of the random choices of the algorithm in the first t rounds is exactly as specified by u .

We say that a vertex $u \in \mathcal{H}$ is a *leaf* if it has no children, i.e., the algorithm terminates given the history of u . In particular, every leaf corresponds to a complete allocation. Let \mathcal{L} be the set of all leaves of the execution tree, and let \mathcal{L}_u denote the set of leaves in the subtree rooted at u .

For each node u , we denote by (X_1^u, \dots, X_n^u) the partial allocation when the algorithm is in state u . For each j , we split the bundle X_j^u into two parts Y_j^u and g^u in the following way which depends on whether j was given an item in line 9 or not. If the last time that j 's bundle changes in the history of u is due to an operation in line 9, i.e., j gets an item because j is an unenvied agent, then we let g^u be the item that was given to j during that operation and $Y_j^u = X_j^u - g^u$. Otherwise, we let $Y_j^u = X_j^u$ and $g^u = \perp$. Note that in both cases it holds that $X_j^u = Y_j^u + g^u$.

Let us first combine the observations made during the analysis of the randomization of the first and the second phase of the algorithm into the following useful lemma.

Algorithm 1 Fair Envy Cycles.

- 1: **Input** A set $[m]$ of items, a set $[n]$ of agents, and a profile $(v_i)_{i \in [n]}$ of valuation functions.
 2: **Output** A complete allocation (X_1, \dots, X_n) that is ex-ante $\frac{1}{2}$ -EF, ex-post $\frac{1}{2}$ -EFX and EF1.

First Phase

- 3: $Z \leftarrow \text{ONE-STEP-PS}(m, n, (v_i)_{i \in [n]})$
 4: $(X^k, p^k)_{k \in [q]} \leftarrow \text{BIRKHOFF-ROUNDING}(Z)$
 5: $X \leftarrow$ an integral allocation sampled from $(X^k, p^k)_{k \in [q]}$

Second Phase

- 6: **while** there is an unallocated item **do**
 7: $x \leftarrow$ an arbitrary unallocated item
 8: **if** there is an unenvied agent **then**
 9: $i \leftarrow$ an arbitrary unenvied agent
 10: $X_i \leftarrow X_i + x$
 11: **else**
 12: $G \leftarrow$ an arbitrary strongly connected component of
 13: the envy graph, without any incoming edges
 14: $c \leftarrow$ a cycle sampled from the distribution obtained
 15: from Lemma 3.12 applied to G
 16: reallocate the bundles along c
 17: **end if**
 18: **end while**
- 19: **return** (X_1, \dots, X_n)

LEMMA 3.15. Consider a node $u \in \mathcal{H}$ of the execution tree, and let u' be the random child of u obtained after performing the next step of the algorithm. Then,

$$v_i(X_i^{u'}) \geq_{\text{SD}} v_i(Y_j^{u'}) \cdot \mathbb{1}[X_j^u \neq X_j^{u'}].$$

PROOF. If u is the root of the tree, the result follows from Lemma 3.8. Otherwise, given the history of u , the algorithm either performs an envy cycle elimination step, in which case the result follows from Lemma 3.13, or the algorithm assigns an unallocated item to an unenvied agent, in which case either j is unenvied or j 's bundle does not change, and the result follows. \square

The proof of the ex-ante fairness of the main algorithm relies on a new notion of *stochastic coverage* which generalizes the notion of stochastic dominance. In order to show the desired inequality $2 \cdot \mathbb{E}[v_i(X_i)] \geq \mathbb{E}[v_i(X_j)]$, we first show that $2 \cdot \mathbb{P}[v_i(X_i) \geq t] \geq \mathbb{P}[v_i(Y_j) \geq t] + \mathbb{P}[v_i(g_j) \geq t]$ which is neatly captured in terms of stochastic coverage as $(v_i(X_i), v_i(X_i)) \geq_{\text{SC}} (v_i(Y_j), v_i(g_j))$. See the following definition.

Definition 3.16 (Stochastic coverage). Given two collections $U = (x_i)_{i \in [r]}$ and $V = (y_j)_{j \in [s]}$ of non-negative random variables (not necessarily independent) for some $r, s \geq 1$, we say that U stochastically covers V , and we write $U \geq_{\text{SC}} V$, if

$$\sum_{i \in [r]} \mathbb{P}[x_i \geq t] \geq \sum_{j \in [s]} \mathbb{P}[y_j \geq t] \quad \text{for all } t > 0$$

This definition captures the standard notion of stochastic dominance between random variables by taking $r = s = 1$. Moreover, note that the definition requires t to be strictly positive, which means,

for instance, that any (even empty) collection of random variables covers any other collection of random variables that are always equal to 0.

REMARK 1. *An alternative way to think about stochastic coverage is in terms of a “credit scheme”. In terms of stochastic coverage, our goal is to show that $(v_i(X_i), v_i(X_i)) \succeq_{\text{SC}} (v_i(Y_j), v_i(g_j))$. The given random variables are defined on the probability space where every outcome is associated with a leaf in the execution tree. Imagine that initially each leaf $u \in \mathcal{L}$ holds two divisible tokens of mass $\mathbb{P}[u]$ each and that it can use some fraction of these tokens to compensate another leaf $v \in \mathcal{L}$ if either $v_i(X_i^u) \geq v_i(Y_j^v)$ or $v_i(X_i^u) \geq v_i(g_j^v)$. The statement $(v_i(X_i), v_i(X_i)) \succeq_{\text{SC}} (v_i(Y_j), v_i(g_j))$ is then equivalent to the existence of a credit scheme where each leaf $v \in \mathcal{L}$ receives a total mass of $\mathbb{P}[v]$ tokens to compensate $v_i(Y_j^v)$ and a total mass of $\mathbb{P}[v]$ tokens to compensate $v_i(g_j^v)$. Hence, we essentially show that the leaves where i is assigned a valuable bundle can compensate all the leaves where j is assigned a valuable bundle (from i 's perspective).*

The following lemma (whose proof is deferred to Appendix E) establishes several useful properties of stochastic coverage.

LEMMA 3.17. *Stochastic coverage satisfies the following properties:*

- (i) *transitivity: for all collections U, V, W , if $U \succeq_{\text{SC}} V$ and $V \succeq_{\text{SC}} W$, then $U \succeq_{\text{SC}} W$.*
- (ii) *concatenability: for all collections S, T, U, V , if $S \succeq_{\text{SC}} T$ and $U \succeq_{\text{SC}} V$, then $(S, U) \succeq_{\text{SC}} (T, V)$.*
- (iii) *disjoint additivity: for all non-negative random variables x and y that are positive on disjoint events, it holds that $(x, y) \succeq_{\text{SC}} (x + y)$ and $(x + y) \succeq_{\text{SC}} (x, y)$.*

We also establish a way to extend stochastic coverage if the covering collection contains strictly more random variables than the covered one, which we use in the form of the following lemma. The proof is deferred to Appendix E.

LEMMA 3.18 (EXTENDABILITY). *For all non-negative random variables x, y, z , if*

- (i) *it holds that $(x, x) \succeq_{\text{SC}} y$,*
 - (ii) *there is an event S such that x, y and z are non-zero only if S holds, and*
 - (iii) *there is some constant threshold δ such that whenever S holds, it also holds that $x \geq \delta \geq z$,*
- then it holds that $(x, x) \succeq_{\text{SC}} (y, z)$.*

Next, we prove the following lemma which is the main ingredient in the proof of Theorem 3.1. The proof proceeds by induction on the execution tree, as illustrated by Figure 3.

LEMMA 3.19. *Letting $w \in \mathcal{L}$ denote the random leaf reached by the algorithm, we have*

$$(v_i(X_i^w), v_i(X_i^w)) \succeq_{\text{SC}} (v_i(Y_j^w), v_i(g_j^w)).$$

PROOF. Let $w \in \mathcal{L}$ be the random leaf reached by a random execution of the algorithm. For any node $u \in \mathcal{H}$, we define the following random variables

$$x_u = \mathbb{1}[w \in \mathcal{L}_u] \cdot v_i(X_i^w), \quad z_u = \mathbb{1}[w \in \mathcal{L}_u] \cdot v_i(X_i^u), \quad y_u = \mathbb{1}[w \in \mathcal{L}_u] \cdot v_i(Y_j^w) \cdot \mathbb{1}[X_j^u \neq X_j^w].$$

Note that y_u includes $v_i(Y_j)$ from all the yellow and green leaves in Figure 3 but not from the red ones. Observe that it always holds that $x_u \geq z_u$ by Lemma 3.3.

First, we use induction on the execution tree to show that for every node u it holds that

$$(x_u, x_u) \succeq_{\text{SC}} (y_u, z_u).$$

This statement holds for all the leaves $u \in \mathcal{L}$ since $x_u = z_u$ and $y_u = 0$.

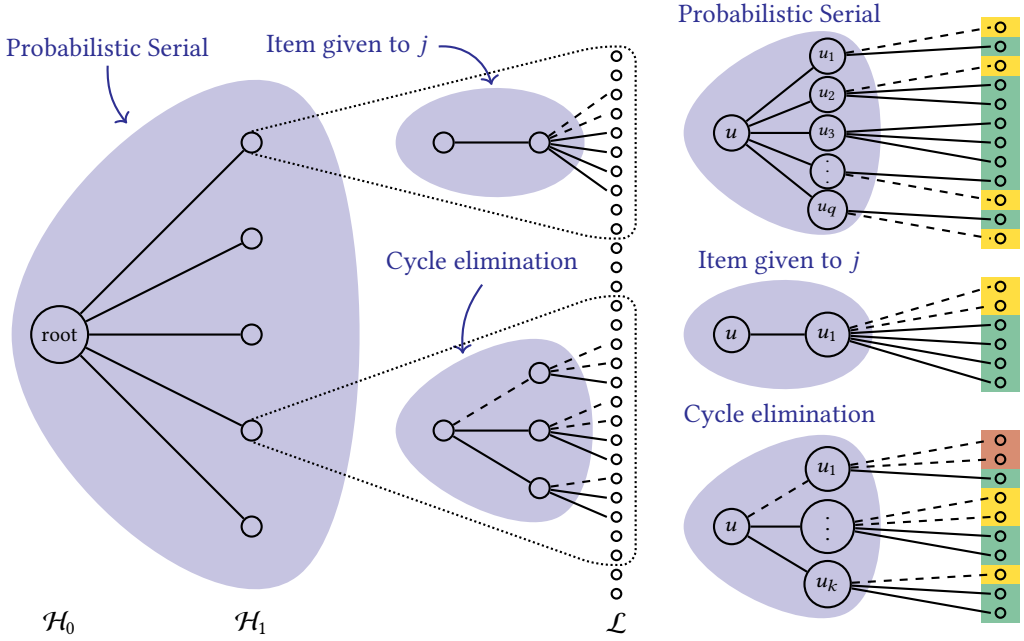


Fig. 3. Illustration of the entire execution tree (left) and its subtrees (right). Paths on which j 's bundle changes are represented by plain edges, and paths on which j 's bundle does not change are represented by dashed edges. The first step is probabilistic serial (top right) and subsequent operations include either assigning an item to an unenvied agent (middle right), or eliminating an envy cycle (bottom right). Let u be some node, and let u_1, \dots, u_k be its children. Consider a leaf w in the subtree rooted at u_t . The leaf w is colored *red* if $X_j^w = X_j^u$, i.e., j 's bundle does not change after reaching u . It is colored *yellow* if $X_j^w = X_j^{u_t}$ and $X_j^{u_t} \neq X_j^u$, i.e., j 's bundle changes only during the first iteration after reaching u . Finally, it is colored *green* if $X_j^w \neq X_j^{u_t}$, i.e., j 's bundle changes during one of the subsequent iterations. Note that a leaf's color is always with respect to some node u , and it can thus be colored differently depending on the chosen node u .

Consider a non-leaf node u and its children u_1, \dots, u_k . By the inductive assumption, it holds that $(x_{u_t}, x_{u_t}) \geq_{SC} (y_{u_t}, z_{u_t})$ for all $1 \leq t \leq k$. This means that $v_i(Y_j)$ from the green leaves in Figure 3 is covered by our induction hypothesis. Observe that

$$\begin{aligned}
 (x_u, x_u) &= (x_{u_1} + \dots + x_{u_k}, x_{u_1} + \dots + x_{u_k}) && \text{(since } x_u = x_{u_1} + \dots + x_{u_k}\text{)} \\
 &\geq_{SC} (x_{u_1}, \dots, x_{u_k}, x_{u_1}, \dots, x_{u_k}) && \text{(by disjoint additivity)} \\
 &\geq_{SC} (y_{u_1}, \dots, y_{u_k}, z_{u_1}, \dots, z_{u_k}) && \text{(by concatenability)} \\
 &\geq_{SC} (y_{u_1} + \dots + y_{u_k}, z_{u_1} + \dots + z_{u_k}) && \text{(by disjoint additivity)}
 \end{aligned}$$

Next, for each child u_t of u , define

$$\mu_{u_t} = \mathbb{1}[w \in \mathcal{L}_{u_t}] \cdot v_i(Y_j^{u_t}) \cdot \mathbb{1}[X_j^u \neq X_j^{u_t} = X_j^w] \quad \text{for } 1 \leq t \leq k$$

so that $y_u = y_{u_1} + \dots + y_{u_k} + \mu_{u_1} + \dots + \mu_{u_k}$. The variables μ_{u_t} include $v_i(Y_j)$ from the yellow leaves in Figure 3, which need to be covered at this induction step. Now, observe that

$$z_{u_1} + \dots + z_{u_k} = \sum_{t=1}^k \mathbb{1}[w \in \mathcal{L}_{u_t}] \cdot v_i(X_j^{u_t}) \quad \text{(by the definition of } z_{u_t}\text{)}$$

$$\begin{aligned}
 &\geq_{\text{SD}} \sum_{t=1}^k \mathbb{1}[w \in \mathcal{L}_{u_t}] \cdot v_i(Y_j^{u_t}) \cdot \mathbb{1}[X_j^u \neq X_j^{u_t}] && \text{(by Lemma 3.15)} \\
 &\geq \sum_{t=1}^k \mathbb{1}[w \in \mathcal{L}_{u_t}] \cdot v_i(Y_j^{u_t}) \cdot \mathbb{1}[X_j^u \neq X_j^{u_t} = X_j^w] \\
 &= \mu_{u_1} + \dots + \mu_{u_k} && \text{(by the definition of } \mu_t)
 \end{aligned}$$

Combining the statements above gives

$$\begin{aligned}
 (x_u, x_u) &\geq_{\text{SC}} (y_{u_1} + \dots + y_{u_k}, z_{u_1} + \dots + z_{u_k}) \\
 &\geq_{\text{SC}} (y_{u_1} + \dots + y_{u_k}, \mu_{u_1} + \dots + \mu_{u_k}) \\
 &\geq_{\text{SC}} y_{u_1} + \dots + y_{u_k} + \mu_{u_1} + \dots + \mu_{u_k} && \text{(by disjoint additivity)} \\
 &= y_u
 \end{aligned}$$

so that $v_i(Y_j)$ from all the green and yellow leaves is covered. To complete the inductive proof, we use the extendability of stochastic coverage (Lemma 3.18) with the event $S = \{w \in \mathcal{L}_u\}$ and the threshold $\delta = v_i(X_j^u)$ to get

$$(x_u, x_u) \geq_{\text{SC}} (y_u, z_u)$$

which completes the induction step.

Now to prove the statement of the lemma, it remains to account for $v_i(g_j^w)$. By disjoint additivity, we get

$$(x_{\text{root}}, x_{\text{root}}) \geq_{\text{SC}} (y_{\text{root}}, z_{\text{root}}) \geq_{\text{SC}} y_{\text{root}}$$

since $z_{\text{root}} = 0$. Note that by strong separation (Lemma 3.10), it holds that $x_{\text{root}} \geq \max_{q \in \mathcal{L}_{\text{root}}} v_i(g_j^q) \geq v_i(g_j^w)$. Hence, we use extendability (Lemma 3.18) once again, now with the full event $S = \{w \in \mathcal{L}_{\text{root}}\}$ and the threshold $\delta = \max_{q \in \mathcal{L}_{\text{root}}} v_i(g_j^q)$, to get

$$(x_{\text{root}}, x_{\text{root}}) \geq_{\text{SC}} (y_{\text{root}}, v_i(g_j^w))$$

which proves the lemma. \square

We are now ready to prove the main theorem.

PROOF OF THEOREM 3.1. The allocation of Algorithm 1 is ex-post $\frac{1}{2}$ -EFX and ex-post EF1 by Lemma 3.4. To show that the outcome is ex-ante $\frac{1}{2}$ -EF, observe that

$$\begin{aligned}
 2 \cdot \mathbb{E}[v_i(X_i)] &= \int_{t \geq 0} \mathbb{P}[v_i(X_i) \geq t] + \mathbb{P}[v_i(X_i) \geq t] \, dt && \text{(by the properties of expectation)} \\
 &\geq \int_{t \geq 0} \mathbb{P}[v_i(Y_i) \geq t] + \mathbb{P}[v_i(g_j) \geq t] \, dt && \text{(by Lemma 3.19)} \\
 &= \mathbb{E}[v_i(Y_j)] + \mathbb{E}[v_i(g_j)] && \text{(by the properties of expectation)} \\
 &= \mathbb{E}[v_i(Y_j) + v_i(g_j)] && \text{(by additivity of expectation)} \\
 &\geq \mathbb{E}[v_i(X_j)] && \text{(by subadditivity of } v_i)
 \end{aligned}$$

which is the desired inequality.

Finally, we prove that Algorithm 1 runs in polynomial time. First, recall that probabilistic serial can be implemented in polynomial time [Bogomolnaia and Moulin, 2001], and that the resulting fractional allocation can be rounded in polynomial time (Theorem 3.6). Second, each iteration of our randomized algorithm can be carried out in polynomial time by Lemma 3.12. Finally, the deterministic envy cycles procedure terminates after polynomially many steps (Lemma 3.5), and

since we only randomize choices that are made arbitrarily in this deterministic procedure, the polynomial running time applies to every outcome of our randomized algorithm. \square

4 Concluding Remarks and Further Directions

Our results significantly advance the state-of-the-art in best-of-both-worlds fairness. We demonstrate that strong ex-ante and ex-post fairness guarantees can be achieved even in settings with subadditive valuations and even with respect to the stronger envy-freeness notion of EFX. Our work suggests several natural problems for future research.

First, while our analysis is tight for our algorithm (see Example D.1), the existence of an allocation that is ex-ante α -EF and ex-post $\frac{1}{2}$ -EFX for any $\alpha > \frac{1}{2}$ remains open.

Second, many of the existing results on EFX rely on the envy cycles procedure, e.g., the existence of a $(\varphi - 1)$ -EFX allocation for any number of additive agents [Amanatidis et al., 2020]. An interesting direction to investigate is whether our techniques can be combined with these results to obtain best-of-both-worlds guarantees for these settings.

Third, while our algorithm produces an execution tree of polynomial depth, and it computes the random allocation in polynomial time, the resulting distribution might have exponential-size support. By Carathéodory's theorem, it is possible to reduce the support size of the final distribution to $n \cdot (n - 1) + 1$. However, it remains a challenge to construct this polynomial-size distribution in polynomial time. Unfortunately, techniques similar to those of Freeman et al. [2020] (who reduce the support size after each iteration) and Aziz [2020] do not apply easily beyond additive valuations.

Fourth, it is possible that further improvements of our bound can be found when considering only ex-post EF1 without EFX. The existence of an ex-ante EF and ex-post EF1 allocation remains an open problem for subadditive valuations. Our investigation did not find any counterexamples for this setting. Although there are several methods to obtain an EF1 allocation for additive valuations, to the best of our knowledge envy-cycles is the only known technique that extends to subadditive valuations. An exciting direction for future work is to examine whether another randomization of envy-cycles could lead to full ex-ante EF. We suspect that this requires finding the right way to randomize the choice of the unenvied agent when assigning the next item. As Example D.1 suggests, the main difficulty in making this improvement is that such randomization cannot simply consider the set of unenvied agents, and must instead take the entire execution tree into account.

Finally, it would be intriguing to explore whether our positive results can be augmented with any type of efficiency guarantees. Prior work shows that random allocations that are ex-ante fractionally-Pareto optimal, ex-ante EF, and ex-post EF1 don't exist even for additive valuations [Freeman et al., 2020]. Similarly, there are instances with additive valuations that admit no allocations that are EFX and Pareto optimal [Plaut and Roughgarden, 2018].

Indeed, the example from [36] shows that ex-post Pareto optimality is incompatible with α -EFX for any $\alpha > 0$ when an instance has 0-valued items. It is plausible that stronger impossibility results than those of Freeman et al. [2020] and Aziz [2020], for weaker notions of efficiency, may apply in our setting. To the best of our knowledge, an approach based on envy cycles is the only known method to obtain the EF1 guarantee for subadditive valuations, and it is well-known that this method does not provide efficiency guarantees. Consequently, positive results in this direction will require the development of new techniques.

Acknowledgments

This project has been partially funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 866132), by an Amazon Research Award, by the NSF-BSF (grant number 2020788), and by a grant from TAU Center for AI and Data Science (TAD).

References

- Hannaneh Akrami, Noga Alon, Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, and Ruta Mehta. 2022. EFX Allocations: Simplifications and Improvements. *CoRR abs/2205.07638* (2022).
- Hannaneh Akrami, Kurt Mehlhorn, Masoud Seddighin, and Golnoosh Shahkarami. 2023. Randomized and Deterministic Maximin-share Approximations for Fractionally Subadditive Valuations. arXiv:2308.14545 [cs.GT]
- Martin Aleksandrov, Haris Aziz, Serge Gaspers, and Toby Walsh. 2015. Online Fair Division: Analysing a Food Bank Problem. In *IJCAI*. AAAI Press, 2540–2546.
- Georgios Amanatidis, Evangelos Markakis, and Apostolos Ntokos. 2020. Multiple Birds with One Stone: Beating 1/2 for EFX and GMMS via Envy Cycle Elimination. In *AAAI*. AAAI Press, 1790–1797.
- Haris Aziz. 2019. *A Probabilistic Approach to Voting, Allocation, Matching, and Coalition Formation*. 45–50.
- Haris Aziz. 2020. Simultaneously Achieving Ex-ante and Ex-post Fairness. In *WINE (Lecture Notes in Computer Science, Vol. 12495)*. Springer, 341–355.
- Haris Aziz, Aditya Ganguly, and Evi Micha. 2023a. Best of Both Worlds Fairness under Entitlements. In *AAMAS*.
- Haris Aziz, Xinhang Lu, Mashbat Suzuki, Jeremy Vollen, and Toby Walsh. 2023b. Best-of-Both-Worlds Fairness in Committee Voting. arXiv:2303.03642 [cs.GT]
- Haris Aziz, Xinhang Lu, Mashbat Suzuki, Jeremy Vollen, and Toby Walsh. 2024. Fair Lotteries for Participatory Budgeting. (2024).
- Haris Aziz and Simon Mackenzie. 2016. A Discrete and Bounded Envy-Free Cake Cutting Protocol for Any Number of Agents. In *FOCS*. IEEE Computer Society, 416–427.
- Moshe Babaioff, Tomer Ezra, and Uriel Feige. 2021. Fair and Truthful Mechanisms for Dichotomous Valuations. In *AAAI*. AAAI Press, 5119–5126.
- Moshe Babaioff, Tomer Ezra, and Uriel Feige. 2022. On Best-of-Both-Worlds Fair-Share Allocations. In *WINE (Lecture Notes in Computer Science, Vol. 13778)*. Springer, 237–255.
- Siddharth Barman, Arpita Biswas, Sanath Krishnamurthy, and Yadati Narahari. 2018. Groupwise maximin fair allocation of indivisible goods. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- Siddharth Barman and Sanath Kumar Krishnamurthy. 2020. Approximation algorithms for maximin fair division. *ACM Transactions on Economics and Computation (TEAC)* 8, 1 (2020), 1–28.
- Ben Berger, Avi Cohen, Michal Feldman, and Amos Fiat. 2022. Almost Full EFX Exists for Four Agents. In *AAAI*. AAAI Press, 4826–4833.
- Umang Bhaskar, A. R. Sricharan, and Rohit Vaish. 2021. On Approximate Envy-Freeness for Indivisible Chores and Mixed Resources. In *APPROX-RANDOM (LIPIcs, Vol. 207)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 1:1–1:23.
- Garrett Birkhoff. 1946. Three observations on linear algebra. *Univ. Nac. Tucuman, Rev. Ser. A* 5 (1946), 147–151.
- Anna Bogomolnaia and Hervé Moulin. 2001. A New Solution to the Random Assignment Problem. *J. Econ. Theory* 100, 2 (2001), 295–328.
- S. Brams and A. Taylor. 1995. An Envy-Free Cake Division Protocol. *The American Mathematical Monthly* 102, 1 (1995), 9–18.
- Eric Budish. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
- Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *EC*. ACM, 305–322.
- Hau Chan, Jing Chen, Bo Li, and Xiaowei Wu. 2019. Maximin-Aware Allocations of Indivisible Goods. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI'19)*. 137–143.
- Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. 2020a. EFX Exists for Three Agents. In *EC*. ACM, 1–19.
- Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, Ruta Mehta, and Pranabendu Misra. 2021. Improving EFX Guarantees through Rainbow Cycle Number. In *EC*. ACM, 310–311.
- Bhaskar Ray Chaudhury, Telikepalli Kavitha, Kurt Mehlhorn, and Alkmini Sgouritsa. 2020b. A Little Charity Guarantees Almost Envy-Freeness. In *SODA*. SIAM, 2658–2672.
- Anton-Hermann Chroust. 1942. Aristotle's conception of equity (epieikeia). *Notre Dame Law Review* 18 (1942), 119–128.
- Uriel Feige, Ariel Sapir, and Laliv Tauber. 2021. A Tight Negative Example for MMS Fair Allocations. In *WINE (Lecture Notes in Computer Science, Vol. 13112)*. Springer, 355–372.
- Michal Feldman, Simon Murras, and Tomasz Ponitka. 2024. On Optimal Tradeoffs between EFX and Nash Welfare. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 9 (2024).
- D. Foley. 1967. Resource allocation and the public sector. *Yale Econ Essays* 7, 1 (1967), 45–98.
- Rupert Freeman, Nisarg Shah, and Rohit Vaish. 2020. Best of Both Worlds: Ex-Ante and Ex-Post Fairness in Resource Allocation. In *EC*. ACM, 21–22.
- Mohammad Ghodsi, Mohammad Taghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. 2018. Fair Allocation of Indivisible Goods: Improvements and Generalizations. In *EC*. ACM, 539–556.

- Pratik Ghosal, Vishwa Prakash H.V., Prajakta Nimbhorkar, and Nithin Varma. 2023. EFX Exists for Four Agents with Three Types of Valuations. *arXiv preprint arXiv:2301.10632* (2023).
- Daniel Halpern, Ariel D. Procaccia, Alexandros Psomas, and Nisarg Shah. 2020. Fair Division with Binary Valuations: One Rule to Rule Them All. In *WINE (Lecture Notes in Computer Science, Vol. 12495)*. Springer, 370–383.
- Martin Hoefer, Marco Schmalhofer, and Giovanna Varricchio. 2023. Best of Both Worlds: Agents with Entitlements. In *AAMAS*.
- Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *EC. ACM*, 125–131.
- J. MacQueen. 1981. Circuit Processes. *The Annals of Probability* 9, 4 (1981), 604 – 610.
- Ryoga Mahara. 2021. Extension of Additive Valuations to General Valuations on the Existence of EFX. In *29th Annual European Symposium on Algorithms (ESA 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- James R Norris. 1998. *Markov chains*. Number 2. Cambridge university press.
- Benjamin Plaut and Tim Roughgarden. 2018. Almost Envy-Freeness with General Valuations. In *SODA. SIAM*, 2584–2603.
- Ariel D. Procaccia. 2020. An answer to fair division’s most enigmatic question: technical perspective. *Commun. ACM* 63, 4 (2020), 118.
- Ariel D. Procaccia and Junxing Wang. 2014. Fair enough: guaranteeing approximate maximin shares. In *EC. ACM*, 675–692.
- Hugo Steinhaus. 1948. The problem of fair division. *Econometrica* 16 (1948), 101–104.
- Walter Stromquist. 1980. How to cut a cake fairly. *The American Mathematical Monthly* 87, 8 (1980), 640–644.
- Hal R Varian. 1974. Equity, envy, and efficiency. *Journal of Economic Theory* 9, 1 (1974), 63–91.
- John Von Neumann. 1953. A certain zero-sum two-person game equivalent to the optimal assignment problem. *Contributions to the Theory of Games* 2, 0 (1953), 5–12.

A The Case of Two Agents

In this section, we provide tradeoffs between ex-post EFX and ex-ante EF for the case of two agents. Our results are demonstrated in Figure 1 (left).

In our algorithms and analysis we use the notion of a β -EFX partition, defined as follows.

Definition A.1 (β -EFX partition). A partition of the items $X = (X_1, X_2)$ is a β -EFX partition for agent i if $v_i(X_1) \geq \beta \cdot v_i(X_2 - g)$ for all $g \in X_2$ and $v_i(X_2) \geq \beta \cdot v_i(X_1 - g)$ for all $g \in X_1$.

Clearly, given a β -EFX partition (X_1, X_2) for agent i , we can allocate either X_1 or X_2 to agent i and preserve β -EFX with respect to agent i .

A.1 Additive Valuations

First, we provide a complete analysis of the case of two agents with additive valuations. Our analysis uses the following technical lemma.

LEMMA A.2 (PARTITION FOR ADDITIVE). *Assume that the valuations are additive. For any agent $i \in \{1, 2\}$, let $X = (X_1, X_2)$ be a partition of the items with $v_i(X_1) \geq v_i(X_2)$ that minimizes $|v_i(X_1) - v_i(X_2)|$ and, subject to this, minimizes $|X_1|$. Then, X is an EFX partition for agent i .*

PROOF. Suppose that $v_i(X_1) \geq v_i(X_2)$ and that it holds that $v_i(X_1 - g) > v_i(X_2)$ for some $g \in X_1$, i.e., X is not an EFX partition for agent i . Consider the partition $X' = (X_1 - g, X_2 + g)$. If $v_i(g) = 0$, then this contradicts the assumption since $|X_1 - g| < |X_1|$. Note that

$$v_i(X'_1) - v_i(X'_2) = v_i(X_1 - g) - v_i(X_2 + g) = v_i(X_1) - v_i(X_2) - 2 \cdot v_i(g)$$

and

$$0 < v_i(g) < v_i(X_1) - v_i(X_2).$$

It follows that

$$-v_i(X_1) + v_i(X_2) < v_i(X'_1) - v_i(X'_2) < v_i(X_1) - v_i(X_2),$$

which contradicts the assumption that X is chosen to minimize $|v_i(X_1) - v_i(X_2)|$. \square

The following proposition provides fairness guarantees for two additive valuations.

PROPOSITION A.3. *Every instance with two additive valuations admits a randomized allocation that is ex-ante EF and ex-post EFX.*

PROOF. Let agents 1 and 2 choose EFX partitions $A = (A_1, A_2)$ and $B = (B_1, B_2)$ respectively, as specified in Lemma A.2. Observe that if agent 2 prefers A_2 over A_1 or agent 1 prefers B_2 over B_1 , then there is a deterministic EF allocation.

Assume that both agents prefer A_1 over A_2 and B_1 over B_2 . Consider the following randomized allocation. With probability $1/2$ give A_1 to 2 and A_2 to 1, and with probability $1/2$ give B_1 to 1 and B_2 to 2. Both of the deterministic allocations are EFX by the assumption.

The expected value of agent 1 for the allocation that agent 1 gets is $(1/2) \cdot (v_1(B_1) + v_1(A_2))$. The expected value of agent 1 for the allocation that agent 2 gets is $(1/2) \cdot (v_1(A_1) + v_1(B_2))$. Since A is chosen to minimize $|v_1(A_1) - v_1(A_2)|$, it holds that

$$v_1(A_1) - v_1(A_2) \leq v_1(B_1) - v_1(B_2)$$

and so

$$(1/2) \cdot (v_1(A_1) + v_1(B_2)) \leq (1/2) \cdot (v_1(B_1) + v_1(A_2))$$

which implies that agent 1 is ex-ante EF. The same argument applies to agent 2. \square

A.2 Subadditive Valuations

In this section, we analyze the case of two agents with subadditive valuations. We begin with the following construction of a $\frac{1}{2}$ -EFX partition that we will use to obtain improved tradeoffs for the case of two agents.

LEMMA A.4 (PARTITION FOR SUBADDITIVE). *Assume that the valuations are subadditive. For any agent $i \in \{1, 2\}$, let $X = (X_1, X_2)$ be a partition of items with $v_i(X_1) \geq v_i(X_2)$ that minimizes $|v_i(X_1) - v_i(X_2)|$, and, subject to this, minimizes $|X_1|$. Then, X is a $\frac{1}{2}$ -EFX partition for agent i .*

PROOF. For the purpose of contradiction, assume that there is some $g \in X_1$ so that $v_i(X_1 - g) > 2 \cdot v_i(X_2)$. First, suppose that $v_i(X_2 + g) \leq v_i(X_1 - g)$. Then,

$$v_i(X_2) \leq v_i(X_2 + g) \leq v_i(X_1 - g) \leq v_i(X_1).$$

This implies that

$$v_i(X_1 - g) - v_i(X_2 + g) \leq v_i(X_1) - v_i(X_2),$$

contradicting the assumption since $|X_1 - g| < |X_1|$.

Next, suppose that $v_i(X_2 + g) > v_i(X_1 - g)$. Observe that

$$\begin{aligned} v_i(X_1 - g) - v_i(X_2) &> v_i(X_2) && \text{(since } 2 \cdot v_i(X_2) < v_i(X_1 - g)\text{)} \\ &\geq v_i([m]) - v_i(X_1) && \text{(since } v_i(X_1) + v_i(X_2) \geq v_i([m])\text{)} \\ &\geq v_i(X_2 + g) - v_i(X_1) && \text{(since } v_i(X_2 + g) \leq v_i([m])\text{)} \end{aligned}$$

which implies that

$$\begin{aligned} &v_i(X_2 + g) - v_i(X_1) < v_i(X_1 - g) - v_i(X_2) \\ \iff &v_i(X_2 + g) - v_i(X_1 - g) < v_i(X_1) - v_i(X_2) \\ \iff &|v_i(X_2 + g) - v_i(X_1 - g)| < |v_i(X_1) - v_i(X_2)| \quad \text{(since } v_i(X_2 + g) \geq v_i(X_1 - g)\text{)}, \end{aligned}$$

contradicting the assumption that X is chosen to minimize $|v_i(X_1) - v_i(X_2)|$. \square

Moreover, there exists an instance with submodular (and hence subadditive) valuations such that the above partition is not a β -EFX partition for agent i for any $\beta > 1/2$; i.e., Lemma A.4 is tight.

PROOF OF TIGHTNESS OF LEMMA A.4. To see that the partition specified in Lemma A.4 cannot provide better guarantees, consider the following submodular valuation over items a, b, c .

	\emptyset	a	b	c	ab	ac	bc	abc
v_1	0	$2 - \varepsilon$	$2 - \varepsilon$	1	$2 - \varepsilon$	$3 - \varepsilon$	$3 - \varepsilon$	$3 - \varepsilon$

One can easily verify that for some small $\varepsilon > 0$, the unique partition $X = (X_1, X_2)$ that minimizes $|v_1(X_1) - v_1(X_2)|$ is given by $X_1 = \{a, b\}$ and $X_2 = \{c\}$.

It holds that

$$v_1(X_1 - b) = v_1(a) = 2 - \varepsilon = (2 - \varepsilon) \cdot v_1(c) = v_1(X_2)$$

which means X is not a β -EFX partition for any $\beta > \frac{1}{2-\varepsilon}$. The result follows by talking $\varepsilon \rightarrow 0$. \square

The following propositions provide fairness tradeoffs for two subadditive valuations. We first present an algorithm that finds an allocation that is ex-ante EF and ex-post $\frac{1}{2}$ -EFX.

PROPOSITION A.5. *Every instance with two subadditive valuations admits a randomized allocation that is ex-ante EF and ex-post $\frac{1}{2}$ -EFX.*

PROOF. We may assume without loss of generality, by uniformly scaling the valuation functions, that each agent’s value for the bundle containing all items is exactly 1. Suppose agent 1 and agent 2 choose $\frac{1}{2}$ -EFX partitions $A = (A_1, A_2)$ and $B = (B_1, B_2)$ respectively, as specified in Lemma A.4. If agent 2 prefers A_2 over A_1 or agent 1 prefers B_2 over B_1 then there is a deterministic EF allocation.

For the remaining case, suppose both agents prefer A_1 to A_2 and B_1 to B_2 . Consider the allocation which, with probability $\frac{1}{2}$, gives A_2 to agent 1 and A_1 to agent 2, and with probability $\frac{1}{2}$ gives B_1 to agent 1 and B_2 to agent 2. By construction, both allocations are $\frac{1}{2}$ -EFX for both agents.

Agent 1’s expected value in this random allocation is $(1/2) \cdot (v_1(B_1) + v_1(A_2))$. The expected value that agent 1 has for agent 2’s bundle is $(1/2) \cdot (v_1(A_1) + v_1(B_2))$. Since A is chosen to minimize $|v_1(A_1) - v_1(A_2)|$, we have

$$\begin{aligned} v_1(B_1) - v_1(B_2) &\geq v_1(A_1) - v_1(A_2) \\ \iff (1/2) \cdot (v_1(B_1) + v_1(A_2)) &\geq (1/2) \cdot (v_1(A_1) + v_1(B_2)) \end{aligned}$$

so the allocation is ex-ante EF for agent 1. A similar argument holds for agent 2. \square

We next give an algorithm that finds an allocation that is ex-ante $\frac{2}{3}$ -EF and ex-post EFX.

PROPOSITION A.6. *Every instance with two subadditive valuations admits a randomized allocation that is ex-ante $\frac{2}{3}$ -EF and ex-post EFX.*

PROOF. Once again, we may assume without loss of generality, by uniformly scaling the valuations, that each agent’s value for the bundle containing all items is exactly 1. Suppose agent 1 chooses any EFX partition $A = (A_1, A_2)$ and agent 2 chooses any EFX partition $B = (B_1, B_2)$. We may assume without loss of generality that $v_1(A_1) \geq v_1(A_2)$ and $v_2(B_1) \geq v_2(B_2)$. Again, if agent 2 prefers A_2 over A_1 or agent 1 prefers B_2 over B_1 , there is a deterministic EF allocation.

For the remaining case, suppose both agents prefer A_1 to A_2 and B_1 to B_2 . Consider the random allocation which, with probability $1/2$, gives A_2 to agent 1 and A_1 to agent 2, and with probability $1/2$ gives B_1 to agent 1 and B_2 to agent 2. Both allocations in the support are EFX for both agents.

Agent 1’s expected value in this random allocation is $(1/2) \cdot (v_1(B_1) + v_1(A_2))$. The expected value that agent 1 has for agent 2’s bundle is $(1/2) \cdot (v_1(A_1) + v_1(B_2))$. By assumption, we have $v_1(B_1) \geq v_1(B_2)$, and by subadditivity, $v_1(B_1) + v_1(B_2) \geq v_1(B_1 \cup B_2) = 1$. Consequently, we have $v_1(B_1) \geq 1/2$. Similarly, we have $v_1(A_1) \geq 1/2$.

If $v_1(A_1) > 2 \cdot v_1(A_2)$, then by subadditivity $v_1(A_1) > 1/2$. Additionally, the set A_1 contains only a single item: suppose for a contradiction that A_1 contains at least two items, and let g_1 and g_2 be two items in A_1 . By subadditivity, $v_1(A_1 - g_1) + v_1(A_1 - g_2) \geq v_1(A_1)$, thus at least one of these items (say g_1) is such that $v_1(A_1 - g_1) \geq (1/2) \cdot v_1(A_1) > v_1(A_2)$, violating the assumption that A is an EFX partition for agent 1. Consequently, A_1 contains a single item of value greater than $1/2$. By monotonicity, this item is in B_1 and not B_2 , hence $A_1 \subseteq B_1$ and $B_2 \subseteq A_2$. Thus we have $v_1(B_1) \geq v_1(A_1)$ and $v_1(A_2) \geq v_1(B_2)$, so this random allocation is ex-ante EF for agent 1.

It remains to consider agent 1’s envy in the case where $v_1(A_1) \leq 2 \cdot v_1(A_2)$. We have

$$\frac{v_1(A_2) + v_1(B_1)}{v_1(A_1) + v_1(B_2)} \geq \frac{v_1(A_2) + v_1(B_1)}{v_1(A_1) + v_1(B_1)} \geq \frac{(1/2) \cdot v_1(A_1) + v_1(B_1)}{v_1(A_1) + v_1(B_1)} \geq \frac{(1/2) \cdot v_1(A_1) + (1/2)}{v_1(A_1) + (1/2)}$$

Since $v_1(A_1) \leq 1$ the above ratio is at least $\frac{2}{3}$, so agent 1 is at least ex-ante $\frac{2}{3}$ -EF. A similar argument applies to the bundle values for agent 2. \square

Finally, the following proposition establishes some impossibility results, showing, among other results, that Proposition A.6 is tight.

PROPOSITION A.7. *For any $0.618 \approx \varphi - 1 < \beta \leq 1$, there exists an instance with no randomized allocation that is simultaneously ex-ante α -EF and ex-post β -EFX, where $\alpha = \frac{\beta+1}{\beta^2+2\beta}$.*

PROOF. Consider the following instance with three items, a , b and c . Take $0 < \varepsilon < 1 - 2\beta$.

	\emptyset	a	b	c	ab	ac	bc	abc
v_1	0	$1 + \varepsilon$	β	β	$1 + \varepsilon$	$1 + \varepsilon$	2β	2β
v_2	0	β	$1 + \varepsilon$	$1 + \varepsilon$	$1 + \beta$	$1 + \beta$	$1 + \varepsilon$	$1 + \beta$

It can be verified that v_1 and v_2 are subadditive (but observe that v_1 is not submodular because $v_1(abc) - v_1(ab) > v_1(ac) - v_1(a)$). There are three deterministic β -EFX allocations: $X^1 = (a, bc)$, $X^2 = (ab, c)$ and $X^3 = (ac, b)$. Let the corresponding probabilities be p_1, p_2 and p_3 . We have

$$\begin{aligned} \mathbb{E}[v_1(X_1)] &= p_1 \cdot (1 + \varepsilon) + p_2 \cdot (1 + \varepsilon) + p_3 \cdot (1 + \varepsilon) = 1 + \varepsilon \\ \mathbb{E}[v_2(X_2)] &= p_1 \cdot (1 + \varepsilon) + p_2 \cdot (1 + \varepsilon) + p_3 \cdot (1 + \varepsilon) = 1 + \varepsilon \\ \mathbb{E}[v_1(X_2)] &= p_1 \cdot 2\beta + p_2 \cdot \beta + p_3 \cdot \beta = \beta \cdot (1 + p_1) \\ \mathbb{E}[v_2(X_1)] &= p_1 \cdot \beta + p_2 \cdot (1 + \beta) + p_3 \cdot (1 + \beta) = 1 + \beta - p_1 \end{aligned}$$

In particular, observe that the maximum envy is minimized when $\beta \cdot (1 + p_1) = 1 + \beta - p_1$, that is, when $p_1 = 1/(\beta + 1)$, in which case we have $\mathbb{E}[v_1(X_2)] = \mathbb{E}[v_2(X_1)] = 1/\alpha$, with $\alpha = \frac{\beta+1}{\beta^2+2\beta}$. \square

As a corollary, we get: (i) there is no random allocation that is ex-ante α -EF and ex-post EFX for $\alpha > \frac{2}{3}$, and (ii) there is no random allocation that is ex-ante EF and ex-post β -EFX for $\beta > \varphi - 1$.

Next, we show that Proposition A.6 is tight even with respect to two *submodular* agents.

PROPOSITION A.8. *For any $\alpha > 2/3$, there is an instance with two submodular valuations that admits no randomized allocation that is ex-ante α -EF and ex-post EFX.*

PROOF. Consider the following instance with three items, a, b, c , for some $\varepsilon > 0$.

	\emptyset	a	b	c	ab	ac	bc	abc
v_1	0	$1/2 + \varepsilon$	$1/2$	$1/2$	1	$1/2 + \varepsilon$	1	1
v_2	0	$1/2$	$1/2 + \varepsilon$	$1/2$	1	1	$1/2 + \varepsilon$	1

It can be verified that v_1 and v_2 are submodular.

Observe that there are only two deterministic EFX allocations, namely $X^1 = (a, bc)$ and $X^2 = (ac, b)$. Consider a lottery X which returns X^1 with probability p and X^2 with probability $1 - p$. We compute

$$\begin{aligned} \mathbb{E}[v_1(X_1)] &= p \cdot v_1(X_1^1) + (1 - p) \cdot v_1(X_1^2) \\ &= p \cdot v_1(a) + (1 - p) \cdot v_1(ac) \\ &= p \cdot (1/2 + \varepsilon) + (1 - p) \cdot (1/2 + \varepsilon) \\ &= 1/2 + \varepsilon \\ \mathbb{E}[v_1(X_2)] &= p \cdot v_1(X_2^1) + (1 - p) \cdot v_1(X_2^2) \\ &= p \cdot v_1(bc) + (1 - p) \cdot v_1(b) \\ &= p \cdot 1 + (1 - p) \cdot (1/2) \end{aligned}$$

and so by symmetry,

$$\begin{aligned} \mathbb{E}[v_2(X_1)] &= p \cdot (1/2) + (1 - p) \cdot 1 \\ \mathbb{E}[v_2(X_2)] &= 1/2 + \varepsilon. \end{aligned}$$

Note that

$$\max(\mathbb{E}[v_1(X_2)], \mathbb{E}[v_2(X_1)]) = \max(p \cdot 1 + (1 - p) \cdot (1/2), p \cdot (1/2) + (1 - p) \cdot (1/2)) \geq 3/4$$

and so at least one agent is not α -EF for any $\alpha > \frac{1/2+\varepsilon}{3/4}$. The result follows by taking $\varepsilon \rightarrow 0$. \square

A.3 General Monotone Valuations

Finally, we complete our analysis of the two-agent case with general (monotone) valuations, for which we show that obtaining any approximate fairness guarantee is impossible. This is cast in the following proposition.

PROPOSITION A.9. *There exists an instance with two agents with monotone valuations that admits no randomized allocation that is ex-ante α -EF and ex-post β -EFX, for any $\alpha > 0$ and $\beta > 0$.*

PROOF. Consider the following instance with three items, a , b , and c . Let $K > 0$.

	\emptyset	a	b	c	ab	ac	bc	abc
v_1	0	1	0	0	1	1	K	K
v_2	0	0	1	0	1	K	1	K

It can be verified that v_1 and v_2 are monotone.

Observe that there are only two deterministic β -EFX allocations (for any $\beta > 0$), namely $X^1 = (a, bc)$ and $X^2 = (ac, b)$. Consider a lottery X which returns X^1 with probability p and X^2 with probability $1 - p$. We compute

$$\begin{aligned}
 \mathbb{E}[v_1(X_1)] &= p \cdot v_1(X_1^1) + (1 - p) \cdot v_1(X_1^2) \\
 &= p \cdot v_1(a) + (1 - p) \cdot v_1(ac) \\
 &= p \cdot 1 + (1 - p) \cdot 1 \\
 &= 1 \\
 \mathbb{E}[v_1(X_2)] &= p \cdot v_1(X_2^1) + (1 - p) \cdot v_1(X_2^2) \\
 &= p \cdot v_1(bc) + (1 - p) \cdot v_1(b) \\
 &= p \cdot K + (1 - p) \cdot 0
 \end{aligned}$$

and so by symmetry,

$$\begin{aligned}
 \mathbb{E}[v_2(X_1)] &= p \cdot 0 + (1 - p) \cdot K \\
 \mathbb{E}[v_2(X_2)] &= 1.
 \end{aligned}$$

Note that

$$\max(\mathbb{E}[v_1(X_2)], \mathbb{E}[v_2(X_1)]) = \max(p \cdot K, (1 - p) \cdot K) \geq K/2$$

and so at least one of the agents is not α -EF for any $\alpha > 2/K$. The result follows by taking $K \rightarrow \infty$. \square

B The Deterministic Envy Cycles Procedure

In this section, we expand the description of the deterministic envy cycles procedure (Algorithm 2).

LEMMA 3.3. *For any agent i , the value of i for her own bundle weakly increases at each step of the envy cycles procedure.*

PROOF. In each step of the second phase, agent i 's bundle can only be modified in one of the following two ways: either agent i is assigned a new item, in which case agent i 's value for her bundle weakly increases by monotonicity, or agent i is involved in a cycle elimination step, in which case agent i gets a bundle that she envied before the cycle elimination step. \square

Algorithm 2 Deterministic Envy Cycles.

-
- 1: **Input** A set $[m]$ of items, a set $[n]$ of agents, and a profile $(v_i)_{i \in [n]}$ of valuation functions.
 - 2: **Output** A complete allocation (X_1, \dots, X_n) that is $\frac{1}{2}$ -EFX and EF1.
-

First Phase

- 3: $(X_1, \dots, X_n) \leftarrow$ an arbitrary weakly separated (Def. 3.2) allocation with $|X_i| = 1$ for all i
-

Second Phase

- 4: **while** there is an unallocated item **do**
 - 5: $x \leftarrow$ an arbitrary unallocated item
 - 6: **if** there is an unenvied agent **then**
 - 7: $i \leftarrow$ an arbitrary unenvied agent
 - 8: $X_i \leftarrow X_i + x$
 - 9: **else**
 - 10: $i_1, \dots, i_k \leftarrow$ an arbitrary directed cycle in the envy graph
 - 11: $(Y_1, \dots, Y_n) \leftarrow Y_{i_w} = X_{i_{w+1}}$ for all $1 \leq w < k$
 - 12: $Y_{i_k} = X_1$
 - 13: $Y_i = X_i$ for $i \notin \{i_1, \dots, i_k\}$
 - 14: $(X_1, \dots, X_n) \leftarrow (Y_1, \dots, Y_n)$
 - 15: **end if**
 - 16: **end while**
-

- 17: **return** (X_1, \dots, X_n)
-

LEMMA 3.4 ([PLAUT AND ROUGHGARDEN, 2018]). *The allocation returned by Algorithm 2 is $\frac{1}{2}$ -EFX and EF1.*

PROOF. Since any matching gives one item to each agent, the partial allocation at the end of the first phase is EFX (which is a stronger property than both $\frac{1}{2}$ -EFX and EF1).

We first show that assigning an unallocated item g to an unenvied agent j does not violate $\frac{1}{2}$ -EFX or EF1. Indeed, since j is unenvied, it holds that $v_i(X_i) \geq v_i(X_j)$ for any other agent i . Additionally, by weak separation, we have $v_i(X_i) \geq v_i(g)$. Hence, for any item $h \in X_i + g - h$, it holds that

$$\begin{aligned}
 v_i(X_i) &\geq (1/2) \cdot (v_i(X_i) + v_i(g)) && \text{(by the above)} \\
 &\geq (1/2) \cdot v_i(X_i + g) && \text{(by subadditivity)} \\
 &\geq (1/2) \cdot v_i(X_i + g - h) && \text{(by monotonicity)}
 \end{aligned}$$

which means that $\frac{1}{2}$ -EFX still holds, and

$$v_i(X_i) \geq v_i(X_j) = v_i((X_j + g) - g)$$

which means that EF1 still holds.

We now show that eliminating an envy cycle preserves EF1 and $\frac{1}{2}$ -EFX. Indeed, after this step, the value of any agent for her own bundle can only increase, and the set of all assigned bundles remains the same. The result follows. \square

LEMMA 3.5. *Algorithm 2 terminates after polynomially many steps.*

PROOF. To prove this, we need to bound the number of steps in the second phase. Recall that in each step, the algorithm either (1) assigns an unallocated item to an unenvied agent, or (2)

eliminates an envy cycle. Note that the number of steps where operation (1) is executed is at most m since the number of unallocated items strictly decreases with every operation (1).

We now show that after every n^2 consecutive steps where operation (2) is executed, there must be a step where operation (1) is executed. It then follows that the number of steps where operation (2) is executed is at most n^2m . To prove the claim, consider a step that eliminates a cycle C in the envy graph. As a result of this step, the arcs in the cycle C disappear from the envy graph. Moreover, every other arc not in C either disappears or is shifted over by one agent in the envy graph. Thus the number of arcs in the envy graph strictly decreases with each cycle elimination step. This proves the claim since the number of arcs in the envy graph is less than n^2 . \square

C The Probabilistic Serial Lottery

In this section we describe the Probabilistic Serial procedure by analyzing the PS-Lottery algorithm of Aziz [2020]. We show that while the algorithms of Aziz [2020] and Freeman et al. [2020] produce outcomes satisfying ex-post EF1, neither of them provides any ex-post EFX guarantees. The following running example is used to illustrate these points.

Example C.1. Consider the following example, with 2 agents having additive valuations over four items, as described in the table below.

	apple	banana	celery	durian
v_1	10	3ϵ	2ϵ	0
v_2	10	2ϵ	3ϵ	0

The PS-Lottery algorithm works as follows. First, make the number of agents and items equal, by possibly adding some dummy items (of zero value), and then creating $k = m/n$ copies of each agent. In Example C.1, the required number of copies is $k = 2$ copies per agent. Denote by 1_1 and 1_2 the two copies of agent 1, and by 2_1 and 2_2 the two copies of agent 2.

Second, simulate the eating procedure of Bogomolnaia and Moulin [2001] for k units of time. Initially, all items are unconsumed, and during the eating process, items get fractionally consumed by agents. For every $1 \leq t \leq k$, during the t -th unit of time, the t -th copy of each agent participates in the eating process and consumes her favourite available item at a constant rate of one item per one unit of time. Notably, multiple agents may be consuming the same item simultaneously. If at any point, an item is fully consumed, then each one of the agents who are currently eating this item switches to her respective next-favourite item. Here, if an agent is indifferent between multiple items, she selects an arbitrary one. Note that after k units of time, every item is fully consumed.

Figure 4 illustrates the eating procedure described above by showing (a) the intermediate state after one unit of time, and (b) the final outcome after two units of time. The outcome can be thought of as a *fractional* allocation of the items to the copied agents, i.e., a real matrix $(Z_{i,r,j})_{i \in [n], r \in [k], j \in [m]}$ where the r -th copy of agent i receives a $Z_{i,r,j}$ fraction of item j .

The PS-Lottery algorithm then decomposes the fractional allocation into a lottery over *integral* allocations using the Birkhoff-von Neumann theorem (see Section 3.2). The resulting lottery for the running example is given below:

$$\begin{matrix} & a & b & c & d & & a & b & c & d & & a & b & c & d \\
 \begin{matrix} 1_1 \\ 1_2 \\ 2_1 \\ 2_2 \end{matrix} & \begin{pmatrix} 0.5 & 0.5 & 0 & 0 \\ 0 & 0.5 & 0 & 0.5 \\ 0.5 & 0 & 0.5 & 0 \\ 0 & 0 & 0.5 & 0.5 \end{pmatrix} & = & 0.5 \cdot & \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} & + & 0.5 \cdot & \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} .
 \end{matrix}$$

images: flaticon.com

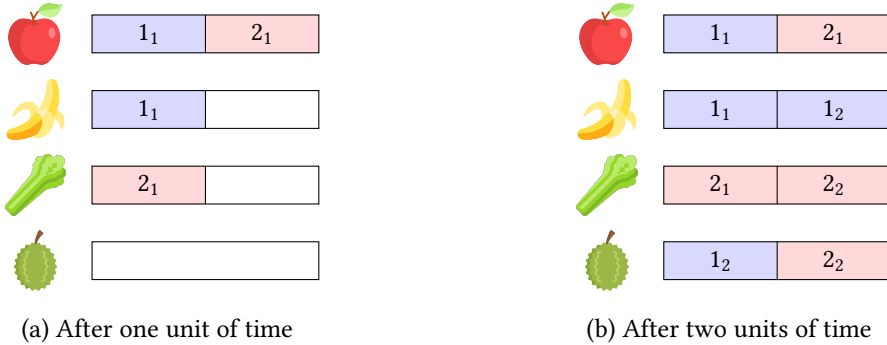


Fig. 4. The eating procedure of Probabilistic Serial, applied to the instance of Example C.1.

Here, the fractional allocation obtained via the eating procedure is given on the left-hand side, and it can be implemented by a lottery over the two integral allocations given in the middle and right-hand side matrices, each with probability $1/2$.

The algorithm then samples one integral allocation according to the obtained probability distribution over integral allocations, and combines the items assigned to all copies of agent i into one bundle that is assigned to agent i in the final allocation. Hence, in the running example, the final allocation is either the allocation where agent 1 receives the bundle $\{a, b\}$ and agent 2 receives the bundle $\{c, d\}$, or the allocation where agent 1 receives the bundle $\{b, d\}$ and agent 2 receives the bundle $\{a, c\}$, each with equal probability.

The result of this process is ex-ante EF by the properties of the eating procedure [Bogomolnaia and Moulin, 2001]. Additionally, Aziz [2020] showed that the outcome is ex-post EF1. However, neither the algorithm of Aziz [2020] (described above) nor the algorithm of Freeman et al. [2020] provides any EFX guarantees, as they always return *balanced* allocations, where all agents receive the same number of items. In general, there may be no balanced allocations that satisfy β -EFX for any $\beta > 0$. This is demonstrated in our running example, where in any balanced allocation, the agent who gets the apple also gets another item, which inevitably violates the β -EFX condition for sufficiently small ε .

D Omitted Propositions and Examples

D.1 The Analysis of Our Main Algorithm is Tight

In this section, we show that the analysis of our main algorithm is tight. We also show that the simple modification of selecting the unenvied agent uniformly at random does not give any better guarantees. We use the following example to demonstrate these facts.

Example D.1. Let $n, k, \varepsilon > 0$. Consider an instance with n agents and $n + k$ items, with the following additive valuations:

	1	2	...	$n - 1$	n	$n + 1$...	$n + k$
v_1	$1 + n\varepsilon$	$1 + (n - 1)\varepsilon$...	$1 + 2\varepsilon$	$1 + \varepsilon$	1	...	1
v_2	$1 + n\varepsilon$	$1 + (n - 1)\varepsilon$...	$1 + 2\varepsilon$	$1 + \varepsilon$	1	...	1
\vdots	\vdots	\vdots	\ddots	\vdots	\vdots	\vdots	\ddots	\vdots
v_{n-1}	$1 + n\varepsilon$	$1 + (n - 1)\varepsilon$...	$1 + 2\varepsilon$	$1 + \varepsilon$	1	...	1
v_n	1	1	...	1	$1 + \varepsilon$	1	...	1

During the eating procedure, each of the agents $1, \dots, n - 1$ consumes a $\frac{1}{n-1}$ fraction of each of the items $1, \dots, n - 1$, and agent n consumes the entire item n . Therefore, after the first phase of the algorithm, there are exactly two agents that no one envies: some agent $i \in \{1, \dots, n - 1\}$ (who gets item $n - 1$) and agent n (who gets item n).

Consider first the case with $k = 1$ to see that no deterministic tie-breaking rule for selecting the unenvied agent in line 9 gives α -EF for any $\alpha > 1/2$. Indeed, no matter which agent gets item $n + 1$ during the second phase of the algorithm, any other agent will envy her with a factor arbitrarily close to 2 for sufficiently small ε .

We also show that selecting the unenvied agent uniformly at random from the set of unenvied agents does not give α -EF for any $\alpha > 1/2$. Here, let us consider $n > k > 1$. Observe that agent n is assigned one of the items $n + 1, \dots, n + k$ with probability $1 - 1/2^k$. Indeed, for every outcome of the first phase of the algorithm, the only outcome of the second phase in which n does not get any of the additional items is where for each $1 \leq w \leq k$, the agent who got item $n - w$ in the first phase of the algorithm, gets item $n + w$ in the second phase, which happens with probability $1/2^k$. Therefore, agent 1's expected value for agent n 's bundle is $(1 + \varepsilon) + (1 - 1/2^k)$ while her expected value for her own bundle is $(1 + \frac{n+2}{2}\varepsilon) + \frac{k}{n}$. The claim follows by taking $k = \log n$ and $n \rightarrow \infty$.

D.2 Ex-ante Guarantees of Random Serial Dictatorship

In this section we establish an upper bound of $\frac{1}{\sqrt{2}}$ on the ex-ante EF guarantee of random serial dictatorship. This improves upon the known bound of $\frac{4}{5}$, from Bogomolnaia and Moulin [2001].

PROPOSITION D.2. *Random serial dictatorship is not ex-ante α -EF for any $\alpha > 1/\sqrt{2}$.*

PROOF. Let $\varepsilon > 0$, and consider the following instance with n additive agents and n items, where k is a parameter (to be determined later) and $p = k/n$.

	a	b	c_1	\dots	c_k	d_1	\dots	d_{n-k-2}
v_1	$1 + \varepsilon$	1	0	\dots	0	0	\dots	0
v_2	1	$1 + \varepsilon$	0	\dots	0	0	\dots	0
v_3	1	0	$1 + \varepsilon$	\dots	$1 + \varepsilon$	ε	\dots	ε
\vdots	\vdots	\vdots	\vdots	\ddots	\vdots	\vdots	\ddots	\vdots
v_n	1	0	$1 + \varepsilon$	\dots	$1 + \varepsilon$	ε	\dots	ε

Consider serial dictatorship with a uniformly random order. Observe that agents $3, \dots, n$ are identical, and let us denote by t the $(k + 1)$ st of them to arrive. We proceed by a case analysis, based on the relative arrival order of agents 1, 2 and t . For each one of these events, we write its probability, and the items chosen by agent 1 and by agent 2.

Relative ordering	Probability	Item chosen by 1	Item chosen by 2
$1, 2, t$	$\frac{1}{2} \frac{k+1}{n-1} \frac{k+2}{n} \approx \frac{p^2}{2}$	a	b
$2, 1, t$	$\frac{1}{2} \frac{k+1}{n-1} \frac{k+2}{n} \approx \frac{p^2}{2}$	a	b
$1, t, 2$	$\frac{k+1}{n-1} \frac{n-k-2}{n} \approx p(1-p)$	a	b
$2, t, 1$	$\frac{k+1}{n-1} \frac{n-k-2}{n} \approx p(1-p)$	some d_i	b
$t, 1, 2$	$\frac{1}{2} \frac{n-k-2}{n-1} \frac{n-k-1}{n} \approx \frac{(1-p)^2}{2}$	b	some d_i
$t, 2, 1$	$\frac{1}{2} \frac{n-k-2}{n-1} \frac{n-k-1}{n} \approx \frac{(1-p)^2}{2}$	some d_i	b

We observe that agent 2 is more likely to receive one of the top items (item a or b) than agent 1, which creates ex-ante envy from agent 1 to agent 2. More precisely

$$\begin{aligned} \mathbb{E}[v_1(X_1)] &= (1 + \varepsilon) \cdot \left(\frac{k+1}{n-1}\right) + 1 \cdot \left(\frac{1}{2} \cdot \frac{n-k-2}{n-1} \cdot \frac{n-k-1}{n}\right) \approx \frac{1+p^2}{2} + \varepsilon p \\ \mathbb{E}[v_1(X_2)] &= 1 \cdot \left(1 - \frac{1}{2} \cdot \frac{n-k-2}{n-1} \cdot \frac{n-k-1}{n}\right) \approx 1 - \frac{(1-p)^2}{2}. \end{aligned}$$

The ex-ante envy guarantee is upper-bounded by the ratio $\mathbb{E}[v_1(X_1)]/\mathbb{E}[v_1(X_2)]$. We take n arbitrarily large, and ε arbitrarily small, to simplify the expression, which gives

$$\frac{\mathbb{E}[v_1(X_1)]}{\mathbb{E}[v_1(X_2)]} \approx \frac{1+p^2}{2-(1-p)^2}$$

The above ratio reaches a minimum of $1/\sqrt{2}$ when $p = \sqrt{2} - 1$. More precisely, we choose $k = \lfloor n \cdot (\sqrt{2} - 1) \rfloor$, which shows that random serial dictatorship is not ex-ante α -EF for any $\alpha > 1/\sqrt{2}$. \square

E Omitted Proofs from Section 3.4

LEMMA 3.17. *Stochastic coverage satisfies the following properties:*

- (i) *transitivity: for all collections U, V, W , if $U \succeq_{SC} V$ and $V \succeq_{SC} W$, then $U \succeq_{SC} W$.*
- (ii) *concatenability: for all collections S, T, U, V , if $S \succeq_{SC} T$ and $U \succeq_{SC} V$, then $(S, U) \succeq_{SC} (T, V)$.*
- (iii) *disjoint additivity: for all non-negative random variables x and y that are positive on disjoint events, it holds that $(x, y) \succeq_{SC} (x + y)$ and $(x + y) \succeq_{SC} (x, y)$.*

PROOF. Properties (i) and (ii) follow directly from the definition. For property (iii), observe that for any random variables x and y that are positive on disjoint events, we have

$$\mathbb{P}[x + y \geq t] = \mathbb{P}[x + y \geq t \text{ and } x > 0] + \mathbb{P}[x + y \geq t \text{ and } y > 0] = \mathbb{P}[x \geq t] + \mathbb{P}[y \geq t]$$

for all $t > 0$, which concludes the proof. \square

LEMMA 3.18 (EXTENDABILITY). *For all non-negative random variables x, y, z , if*

- (i) *it holds that $(x, x) \succeq_{SC} y$,*
 - (ii) *there is an event S such that x, y and z are non-zero only if S holds, and*
 - (iii) *there is some constant threshold δ such that whenever S holds, it also holds that $x \geq \delta \geq z$,*
- then it holds that $(x, x) \succeq_{SC} (y, z)$.*

PROOF. Assume that all three properties hold. For all $t > \delta$, we have

$$\mathbb{P}[z \geq t] = \mathbb{P}[z \geq t \text{ and } S \text{ holds}] + \mathbb{P}[z \geq t \text{ and } S \text{ does not hold}] = 0$$

and thus $2 \cdot \mathbb{P}[x \geq t] \geq \mathbb{P}[y \geq t] + \mathbb{P}[z \geq t]$. For all $0 < t \leq \delta$, we have

$$2 \cdot \mathbb{P}[x \geq t] = 2 \cdot \mathbb{P}[S \text{ holds}] \geq \mathbb{P}[y > 0] + \mathbb{P}[z > 0] \geq \mathbb{P}[y \geq t] + \mathbb{P}[z \geq t],$$

which concludes the proof. \square