

Methods



Best of Both Worlds: Ex Ante and Ex Post Fairness in Resource Allocation

Haris Aziz,^a Rupert Freeman,^{b,*} Nisarg Shah,^c Rohit Vaish^d

^aSchool of Computer Science and Engineering, University of New South Wales, Sydney, New South Wales 2035, Australia; ^bDarden School of Business, University of Virginia, Charlottesville, Virginia 22903; ^cDepartment of Computer Science, University of Toronto, Toronto, Ontario M5T 3A1, Canada; ^dDepartment of Computer Science and Engineering, Indian Institute of Technology Delhi, New Delhi 110 016, India

*Corresponding author

Contact: haziz@cse.unsw.edu.au,  <https://orcid.org/0000-0001-9757-4956> (HA); freemanr@darden.virginia.edu,

 <https://orcid.org/0000-0003-4744-9449> (RF); nisarg@cs.toronto.edu (NS); rvaish@iitd.ac.in,  <https://orcid.org/0000-0002-4417-3718> (RV)

Received: August 14, 2020

Revised: December 23, 2021;
October 12, 2022

Accepted: December 2, 2022

Published Online in *Articles in Advance*:
January 30, 2023

Area of Review: Decision Analysis

<https://doi.org/10.1287/opre.2022.2432>

Copyright: © 2023 INFORMS

Abstract. We study the problem of allocating indivisible goods among agents with additive valuations. When randomization is allowed, it is possible to achieve compelling notions of fairness such as envy-freeness, which states that no agent should prefer any other agent's allocation to their own. When allocations must be deterministic, achieving exact fairness is impossible but approximate notions such as envy-freeness up to one good can be guaranteed. Our goal in this work is to achieve both simultaneously, by constructing a randomized allocation that is exactly fair ex ante (before the randomness is realized) and approximately fair ex post (after the randomness is realized). The key question we address is whether ex ante envy-freeness can be achieved in combination with ex post envy-freeness up to one good. We settle this positively by designing an efficient algorithm that achieves both properties simultaneously. The algorithm can be viewed as a desirable way to instantiate a lottery for the probabilistic serial rule. If we additionally require economic efficiency, we obtain three impossibility results that show that ex post or ex ante Pareto optimality is impossible to achieve in conjunction with combinations of fairness properties. Hence, we slightly relax our ex post fairness guarantees and present a different algorithm that can be viewed as a fair way to instantiate a lottery for the maximum Nash welfare allocation rule.

Funding: This work was supported by the Office of Naval Research [Grant N00014-17-1-2621] and DST INSPIRE [Grant DST/INSPIRE/04/2020/000107].

Keywords: fair resource allocation • indivisible goods • randomization and approximation

1. Introduction

Allocating resources in nonmonetary environments is commonplace in social institutions. Estates must be divided among beneficiaries, jointly held assets split between partners in a divorce, tasks assigned to employees, and educational resources distributed among public schools. Versions of this resource allocation problem—ranging from applied to theoretical models—have been studied in various research communities including mathematics, economics, operations research, and computer science (Brams and Taylor 1996, Matoušek et al. 2003, Moulin 2003, Brandt et al. 2016). A central concern in these settings is *fairness*: how can these decisions be made to not systematically disadvantage individuals or groups of people?

In this work, we consider a fundamental problem in fair resource allocation. A set of m indivisible goods are to be divided among a set of n agents who have different preferences over the goods. An agent's preferences are expressed through a valuation function

that assigns a value to every subset of the goods: We use the terms “value,” “valuation,” and “utility” interchangeably. We restrict ourselves to the case of additive valuation functions, in which each agent has a value for each good and their value for a subset is simply the sum of their values for the goods in the subset. Additive valuations ignore complementary and substitutive effects that may occur in practice, but they are an appealing tradeoff between simplicity and expressiveness. For instance, the adjusted winner procedure (Brams and Taylor 1996) that is commonly used in dispute resolution assumes additive valuations, as does the popular fair division website Spliddit (<http://spliddit.org>).

A particularly appealing notion of fairness is *envy-freeness* (EF) (Gamow and Stern 1958, Foley 1967), which requires that no agent values the resources allocated to another agent more than the resources allocated to herself. When allocations are deterministic, it is not always possible to achieve envy-freeness; imagine two agents liking a single good, which must be

given to one of them, leaving the other envious. The fundamental unfairness present in deterministic allocations motivates the use of randomized mechanisms instead. Indeed, many mechanisms used in practice exploit lotteries to, for example, determine a priority ordering over the agents. With the power of randomization, it is easy to achieve *ex ante* EF: if a single agent is chosen uniformly at random and then allocated all the goods, no agent envies any other in expectation. However, this allocation induces a large amount of envy *ex post*, because one agent receives everything and all others receive nothing.

Although we cannot eliminate *ex post* envy, a significant body of work has focused on defining and achieving approximate fairness for deterministic allocations (Bouveret et al. 2016). One compelling notion is *EF up to one good* (EF1) (Lipton et al. 2004, Budish 2011), which requires that the envy of any agent toward another agent can be eliminated by the removal of at most one good from the envied agent's allocation. A deterministic EF1 allocation can always be achieved. For instance, it is known that the round-robin method—where the agents choose goods one at a time in a repeating fixed order—is guaranteed to output an EF1 allocation. Agents who come later in the ordering may envy those who come early but only up to a single good. One might hope that if the order of the agents was chosen randomly among the $n!$ possible orderings then *ex ante* EF could additionally be achieved, but this is not the case (see additional discussion in Section 3). Other methods of achieving deterministic EF1 allocation are known (Lipton et al. 2004, Barman et al. 2018, Caragiannis et al. 2019, Brustle et al. 2020) but do not naturally lend themselves to exploiting randomization.

This motivates a natural question. *Can we retain EF up to one good as an ex post guarantee, and simultaneously obtain (exact) EF ex ante?* In other words, can we always randomize over EF1 allocations such that the resulting randomized allocation is EF? We show that the answer to this natural and elegant question is *yes*. More generally, we study various combinations of *ex ante* and *ex post* fairness and efficiency guarantees and identify combinations that can (and cannot) be achieved simultaneously. The efficiency concepts that we examine are all based on the *Pareto optimality* (PO) principle: We want to identify allocations such that there does not exist another allocation that all agents weakly prefer and at least some agent strictly prefers. Our constructive results yield efficient algorithms; these improve on prior algorithms that provide either only *ex ante* or only *ex post* guarantees, thus paving the way for fairer resource allocation in practice.

1.1. Our Results

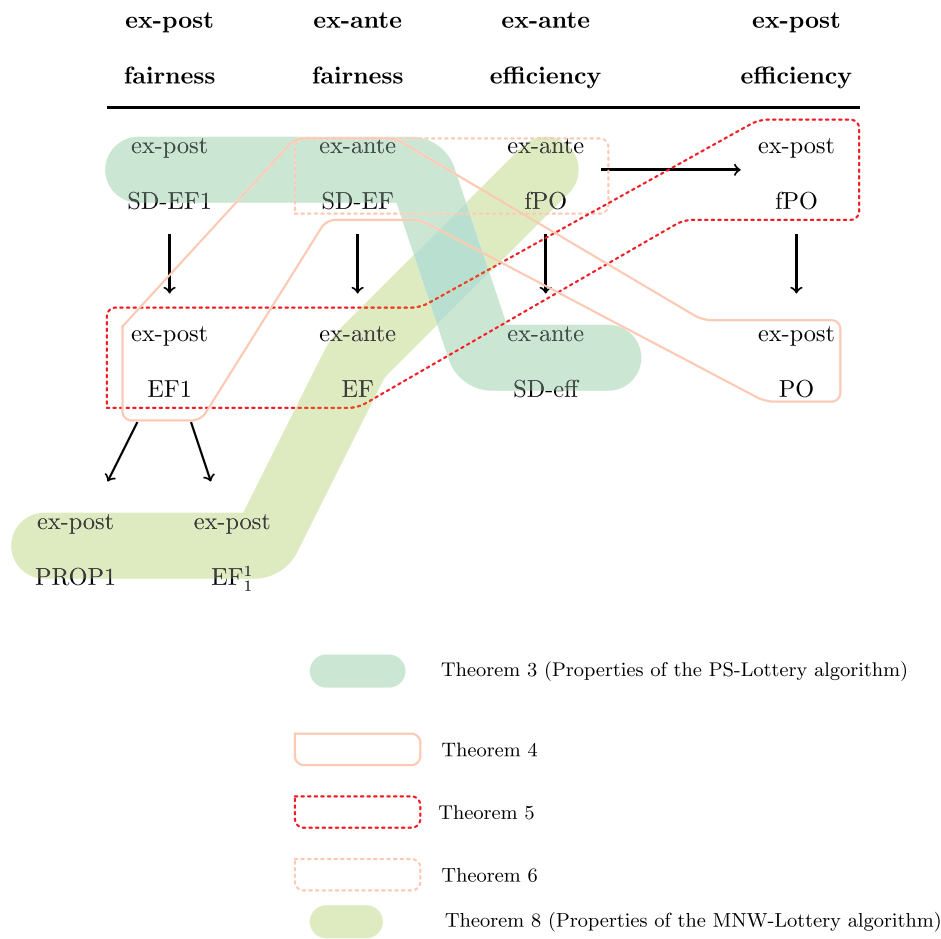
Our first technical result is an algorithm called the Probabilistic Serial-lottery (PS-lottery) algorithm (Algorithm 1) that simultaneously achieves *ex ante* EF and *ex post* EF1

(Theorem 2). The algorithm takes only the agents' ordinal preferences over goods as input and achieves the two guarantees simultaneously with respect to all additive utilities consistent with these ordinal preferences. Following the fair division literature, these stronger guarantees can be phrased as *ex ante* SD-EF and *ex post* SD-EF1, where SD stands for (first-order) stochastic dominance. The algorithm is presented in Section 3. It calls the probabilistic serial algorithm (Bogomolnaia and Moulin 2001) and Birkhoff's decomposition algorithm as subroutines. The algorithm's outcome is *ex ante* equivalent to the outcome of the probabilistic serial rule. In particular, it can be viewed as a desirable way to instantiate a lottery for the *ex ante* outcome of the probabilistic serial rule. We show how the algorithm can be further modified using parametric network flows to additionally achieve both *ex ante* and *ex post* versions of SD-Efficiency (Theorem 3). SD-Efficiency can be viewed as an ordinal and very weak version of PO (with respect to additive valuations). If an allocation is not SD efficient, then there exists another allocation that gives each agent weakly more, and at least one agent strictly more, utility for all utility functions consistent with the ordinal preferences.

If we additionally want to achieve PO, which states that it should be impossible to find an allocation that improves some agent's utility without reducing any other agent's, then the PS-lottery algorithm can be viewed as being maximal in the sense of the following impossibility results that we prove in Section 4 (see Figure 1 for a visual illustration). First, it is impossible to achieve *ex ante* SD-EF, *ex post* EF1, and *ex post* PO (Theorem 4). Second, achieving *ex ante* EF and *ex post* EF1 along with *ex post* fractional PO (a stronger notion of efficiency than *ex post* PO) is also impossible (Theorem 5). Third, *ex ante* fractional PO (i.e., with respect to the randomized allocation) and *ex ante* SD-EF are incompatible (Theorem 6).

In Section 5, we show that strong *ex ante* guarantees—in terms of both fairness and economic efficiency—can be achieved if we are willing to compromise on the *ex post* guarantee. In particular, we are able to achieve *ex ante* *group fairness* (GF) (Conitzer et al. 2019), which generalizes both EF (this is weaker than the SD-EF from Theorem 6) and fractional PO, in conjunction with two *ex post* fairness properties that are incomparable but are both implied by EF1: *proportionality up to one good* (Prop1) (Conitzer et al. 2017) and *EF up to one good more-and-less* (EF₁⁺) (Barman and Krishnamurthy 2019); see Theorem 8 for a formal statement. Our algorithm applies the rounding technique of Budish et al. (2013) to the well-known maximum Nash welfare (MNW) allocation; in particular, it coincides with the *ex ante* MNW outcome. The main technical contribution in this section is to tighten the analysis of Budish et al. (2013) in a way that implies the axiomatic properties that we desire.

Figure 1. Logical Relations Between Fairness and Efficiency Concepts



Notes. An arrow from (a) to (b) denotes that (a) implies (b). The properties in green and lime are simultaneously satisfied by algorithms. The combined properties in the pink or red shapes are impossible to simultaneously satisfy.

Our results indicate that understanding the compatibility of fairness and efficiency from a combined ex ante and ex post perspective provides interesting challenges that can also be explored in other allocation and collective-decision problems.

1.2. Related Work

A large body of work in computer science and economics has focused on finding exactly ex ante fair randomized allocations and approximately fair deterministic allocations, and we cite those works as appropriate throughout the paper. Combining the two approaches was recently listed as an “interesting challenge” by Aziz (2019); however, little work has focused on this problem. Two exceptions are Aleksandrov et al. (2015) and Budish et al. (2013). Aleksandrov et al. (2015) consider randomized allocation mechanisms for an online fair division problem and analyze their ex ante and ex post fairness guarantees. The style of their results is very similar to ours; however, they restrict attention to binary utilities, which simplifies the problem significantly.

Budish et al. (2013) study the problem of implementing a general class of random allocation mechanisms subject to ex post constraints, and we build on this work in Section 5. The ex post constraints that Budish et al. (2013) establish are not the same as ours; in particular, they do not consider ex post axiomatic guarantees from the fair division literature as we do.

In the random assignment literature in economics, the idea of constructing a fractional assignment and implementing it as a lottery over pure assignments was introduced by Hylland and Zeckhauser (1979). Later work has studied both ex ante and ex post fairness and efficiency guarantees provided by mechanisms in this setting (Abdulkadiroğlu and Sönmez 1998, Bogomolnaia and Moulin 2001, Chen and Sönmez 2002, Nesterov 2017), but most of this work studies ordinal utilities and does not consider approximate notions of ex post fairness. (The standard random assignment setting has n agents, n goods, and requires that each agent receive exactly one good. The notions of ex post fairness that we use in this work are vacuous

in this restricted setting.) Gajdos and Tallon (2002) study the relationship between ex ante and ex post fairness but in their model the randomness comes from nature, not the allocation rule. Other works consider the problem of implementing a fractional outcome over deterministic outcomes subject to (possibly soft) constraints (Budish et al. 2013, Akbarpour and Nikzad 2020), but the constraints allowed by these papers do not fully capture our ex post fairness notions.

Although we focus on ex ante EF and ex post EF1 (and relaxations of these properties), many other definitions of fairness have been studied in the literature on resource allocation. For example, *equitability* requires that all agents receive the same utility (Dubins and Spanier 1961, Alon 1987, Cechlárová et al. 2013) and permits additive “up to one good” relaxations (Gourvès et al. 2014, Freeman et al. 2019). Although envy-freeness is often achieved by maximizing the product of utilities (see Section 5), equitability is achieved by maximizing a different common welfare function: the minimum utility. Other fairness notions have been considered in follow-up work to this paper. Halpern et al. (2020) study ex ante and ex post fairness in the context of binary valuations and show that the fractional MNW rule can be implemented as a distribution over deterministic MNW allocations. In a similar vein, Babaioff et al. (2021) pursue our “best of both worlds” approach but consider alternative fairness concepts related to maximin fair share (Budish 2011). In another recent paper, Caragiannis et al. (2021) consider a notion called *interim envy-freeness* that is between the stringent notion of ex post EF and the weak notion of ex ante EF.

2. Preliminaries

For any positive integer $r \in \mathbb{N}$, define $[r] := \{1, \dots, r\}$. Let $N = [n]$ denote a set of *agents*, and M denote a set of *goods* where $m := |M|$.

2.0.1. Fractional and Randomized Allocations. A *fractional allocation* of the goods in M to the agents in N is specified by a nonnegative $n \times m$ matrix $A \in [0, 1]^{n \times m}$ such that for every good $g \in M$, we have $\sum_{i \in N} A_{i,g} = 1$; here, $A_{i,g}$ denotes the fraction of good g assigned to agent i .

A fractional allocation A is *integral* if $A_{i,g} \in \{0, 1\}$ for every $i \in N$ and $g \in M$. For integral allocations, we will find it convenient to denote the binary vector $A_i = (A_{i,g})_{g \in M}$ as a set $A_i := \{g \in M : A_{i,g} = 1\}$. We will refer to A_i as the *bundle* of goods assigned to agent i and denote the allocation A as an ordered tuple of bundles $A = (A_1, \dots, A_n)$. When we simply say “an allocation,” it will mean a fractional allocation, unless otherwise clear from the context. For notational clarity, we use the letters X or Y for fractional allocations and write A

or B for integral allocations. Let \mathcal{X} be the set of all fractional allocations.

A *randomized allocation* is a lottery over integral allocations (we denote them by bold letters for clarity). Formally, a randomized allocation \mathbf{X} is specified by a set of $\ell \in \mathbb{N}$ ordered pairs $\{(p^k, A^k)\}_{k \in [\ell]}$, where, for every $k \in [\ell]$, A^k is an integral allocation implemented with probability $p^k \in [0, 1]$, and $\sum_{k \in [\ell]} p^k = 1$. The *support* of \mathbf{X} is the set of integral allocations $\{A^1, \dots, A^\ell\}$.

A randomized allocation $\mathbf{X} := \{(p^k, A^k)\}_{k \in [\ell]}$ is naturally associated with the fractional allocation $X := \sum_{k \in [\ell]} p^k A^k$, where $X_{i,g}$ is the (marginal) probability of agent i receiving good g under \mathbf{X} . In this case, we say that randomized allocation \mathbf{X} *implements* fractional allocation X . There may be many randomized allocations implementing a given fractional allocation.

2.0.2. Preferences. Each agent $i \in N$ has an *additive* valuation function v_i , where $v_i(g) \geq 0$ denotes the agent’s utility for fully receiving good $g \in M$. Note that v_i induces a weak order \succsim_i over goods where $g \succsim_i g'$ if and only if $v_i(g) \geq v_i(g')$, and $g \succ_i g'$ if and only if $v_i(g) > v_i(g')$. The utility of agent i under an allocation $X \in \mathcal{X}$ is given by, with slight abuse of notation, $v_i(X_i) = \sum_{g \in M} X_{i,g} \cdot v_i(g)$. We assume that for each good $g \in M$, there exists at least one agent $i \in N$ with $v_i(g) > 0$. This is without loss of generality as goods valued zero by everyone can be allocated arbitrarily.

2.0.3. Allocation Rule. A fair division instance I is defined by the triple $(N, M, (v_i)_{i \in N})$. We let \mathcal{I} denote the set of all instances. An *allocation rule* $f : \mathcal{I} \rightarrow 2^{\mathcal{X}}$ maps instances to (sets of) allocations.

Example 1. Consider an instance with two agents 1 and 2 and four goods g_1, g_2, g_3, g_4 . Then, the following is one possible fractional allocation:

$$X = \begin{bmatrix} 1/2 & 1/2 & 1/2 & 1/2 \\ 1/2 & 1/2 & 1/2 & 1/2 \end{bmatrix}.$$

The fractional allocation X can be achieved by the following randomized allocation (other choices are also possible) that is based on a probability distribution over two integral allocations:

$$X = \frac{1}{2} \cdot \begin{bmatrix} 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} + \frac{1}{2} \cdot \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}.$$

We now discuss a number of properties concerning fairness and efficiency of allocations and allocation rules.

2.0.4. Ex Ante and Ex Post Properties. Central to this paper is the distinction between a property holding ex ante and ex post. For any property $\langle P \rangle$ defined for a fractional allocation, we say that a randomized allocation \mathbf{X} satisfies $\langle P \rangle$ *ex ante* if the fractional allocation X it implements satisfies $\langle P \rangle$. Similarly, for any property $\langle Q \rangle$ defined for an integral allocation, we say that a randomized allocation \mathbf{X} satisfies $\langle Q \rangle$ *ex post* if every integral allocation in its support satisfies $\langle Q \rangle$.

2.1. Fairness Properties

Our central fairness concept is envy-freeness, which plays a fundamental role in the economic literature on fairness. Envy-freeness requires that an agent has weakly more value for their own bundle than any other agent's bundle.

Definition 1 (Envy-freeness (EF); Gamow and Stern 1958, Foley 1967). An allocation X is *envy-free* if for every pair of agents $i, j \in N$, we have $v_i(X_i) \geq v_i(X_j)$.

One can also consider a more stringent notion of envy-freeness that requires that an agent has weakly more value for their own bundle than for any other agent's bundle with respect to all additive valuation functions consistent with the agent's ordinal preferences. The requirement can be captured via the first-order SD relation as follows. Given allocations X and Y , we say that agent i *SD-prefers* X_i to Y_i , written $X_i \succsim_i^{SD} Y_i$, if for every good $g \in M$, we have that $\sum_{g' \in M: g' \prec_i g} X_{i,g'} \geq \sum_{g' \in M: g' \prec_i g} Y_{i,g'}$. It is easy to check that $X_i \succsim_i^{SD} Y_i$ is equivalent to $v'_i(X_i) \geq v'_i(Y_i)$ under every additive valuation v'_i consistent with the ordinal preference relation \succsim_i . We write $X_i \succ_i^{SD} Y_i$ if $X_i \succsim_i^{SD} Y_i$ holds but $Y_i \not\prec_i^{SD} X_i$ does not. Based on the SD preference relation, we can define SD-EF as follows.

Definition 2 (SD-envy-freeness (SD-EF); Bogomolnaia and Moulin 2001). An allocation X is *SD-envy-free* if for every pair of agents $i, j \in N$, we have $X_i \succsim_i^{SD} X_j$.

Both envy-freeness and SD-envy-freeness are concepts that apply to both fractional allocations and integral allocations. Hence, each leads to an ex ante and an ex post version for randomized allocations: The former is satisfied if the induced fractional allocation is (SD-)envy-free, and the latter if each integral allocation in the support is (SD-)envy-free.

As we have already seen, an integral allocation satisfying envy-freeness is not guaranteed to exist. In view of this, a literature has developed that focuses on relaxations of envy-freeness, including the following property requiring that pairwise envy can be eliminated by removing a single good from the envied agent's bundle.

Definition 3 (Envy-freeness up to One Good (EF1); Lipton et al. 2004, Budish 2011). An integral allocation

A is *envy-free up to one good* if for every pair of agents $i, j \in N$ such that $A_j \neq \emptyset$, we have $v_i(A_i) \geq v_i(A_j \setminus \{g\})$ for some good $g \in A_j$.

Because EF1 is defined only for integral allocations, a randomized allocation can be ex post EF1 (if each integral allocation in its support satisfies EF1) but ex ante EF1 is not well defined. We also consider the following natural strengthening of EF1.

Definition 4 (SD-envy-freeness up to One Good (SD-EF1)). An integral allocation A is *SD-envy-free up to one good* if for every pair of agents $i, j \in N$ such that $A_j \neq \emptyset$, we have $A_i \succsim_i^{SD} A_j \setminus \{g\}$ for some good $g \in A_j$.

Equivalently, an allocation is SD-EF1 if it is EF1 under any additive valuation functions of the agents consistent with $(\succsim_i)_{i \in N}$.

2.2. Efficiency

Next, we discuss economic efficiency of allocations.

Definition 5 (Fractional PO (fPO) and PO (PO)). An allocation X is *fractionally Pareto optimal* if there is no fractional allocation Y that Pareto-dominates it, that is, satisfies $v_i(Y_i) \geq v_i(X_i)$ for all agents $i \in N$ and at least one inequality is strict. An integral allocation A is *Pareto optimal* if there is no integral allocation B that Pareto-dominates it.

Definition 5 defines one ex ante property (ex ante fPO) and two ex post properties (ex post fPO and ex post PO). For fractional allocations, fPO is traditionally just referred to as PO, so we will slightly abuse terminology by using fPO and PO interchangeably to describe the ex ante property. However, for integral allocations, fPO is stronger than PO (Barman et al. 2018) so it is important to distinguish between the two in the ex post sense.

Proposition 1. *If a randomized allocation is ex ante fractionally Pareto optimal, then it is also ex post fractionally Pareto optimal.*

Proof. If a randomized allocation $\mathbf{X} := \{(p^k, A^k)\}_{k \in [\ell]}$ implementing a fractional allocation X is not ex post fPO, then for some $k \in [\ell]$, the integral allocation A^k must be Pareto dominated by a fractional allocation, say Y . Then, the fractional allocation $X' := p^k \cdot Y + \sum_{r \in [\ell] \setminus \{k\}} p^r \cdot A^r$ Pareto-dominates X , which implies that \mathbf{X} is not ex ante fPO. \square

We can also consider a weak version of efficiency that is based on the SD relation. It requires that there should not exist an alternative allocation that is a Pareto improvement for all additive valuation functions of the agents consistent with $(\succsim_i)_{i \in N}$.

Definition 6 (SD-Efficiency and Weak SD-Efficiency; Bogomolnaia and Moulin 2001). An allocation X is *SD*

efficient if there is no fractional allocation Y such that $Y_i \succsim_i^{SD} X_i$ for all $i \in N$ and $Y_i \succ_i^{SD} X_i$ for some $i \in N$. Additionally, we say that X is *weakly SD efficient* if there is no fractional allocation Y such that $Y_i \succ_i^{SD} X_i$ for all $i \in N$.

SD-Efficiency is also referred to as ordinal efficiency in the literature. Figure 1 illustrates the relations between fairness and efficiency concepts.

3. PS-Lottery Algorithm

This section describes our main result that ex ante envy-freeness can be achieved in conjunction with ex post envy-freeness up to one good.

A natural approach toward this question is to start with the round-robin method. Under this method, we fix an agent ordering, and then agents take turns picking one good at a time in a cyclic fashion. At each step, an agent picks their most valuable good that is still available. It is well known that for any agent ordering, this method produces an integral allocation that is EF1 (Caragiannis et al. 2019). Furthermore, it is easy to see that uniformly randomizing the agent ordering—the so-called *randomized round-robin method*—also achieves a relaxation of ex ante EF called ex ante *proportionality*.

Definition 7 (Proportionality (Prop); Steinhaus 1948). An allocation X is *proportional* if for every agent $i \in N$, we have $v_i(X_i) \geq \frac{1}{n} \cdot v_i(M)$.

However, it can be shown that the randomized round-robin method could fail to achieve ex ante EF; this follows from the observation of Bogomolnaia and Moulin (2001) that the random priority rule (which is randomized round-robin when the number of agents n is equal to the number of goods m) is not ex ante envy-free.

Instead of starting from a method that guarantees ex post EF1 and using it to achieve ex ante EF, let us do the opposite: Start from a fractional EF allocation and implement it using integral EF1 allocations. *Probabilistic serial* is a well-studied algorithm that produces a fractional envy-free allocation (Bogomolnaia and Moulin 2001, Kojima 2009). The algorithm starts with all agents simultaneously eating one of their respective favorite goods at the same constant speed. (Ties between goods can be broken arbitrarily within the algorithm.) Once a good is completely consumed by a subset of agents, each of those agents proceeds to eating one of their favorite available goods at the same speed. The algorithm terminates when all goods have been eaten, and the fraction of each good consumed by an agent is allocated to her. A useful property of this algorithm is that it only uses the ordinal preferences of agents over goods, and computes an allocation that is ex ante envy-free for *any* additive utilities

consistent with the ordinal preferences; thus, it is *SD-envy-free*. Although described as a continuous rule where agents eat infinitesimal amounts, the PS outcome can be computed by a discrete algorithm in polynomial time $O(nm)$ (Kojima 2009).

The challenge that we address here is to find an implementation of the probabilistic serial allocation that additionally satisfies ex post EF1. We do this using an algorithm that we refer to as the PS-lottery algorithm, which only uses the underlying ordinal preferences of the agents. We will make use of the following classical theorem (Birkhoff 1946, von Neumann 1953, Johnson et al. 1960, Plummer and Lovász 1986). A square matrix is called bistochastic if its entries are nonnegative and each of its rows and columns sum to one; a bistochastic matrix is called a permutation matrix if its entries are in $\{0, 1\}$.

Theorem 1 (Birkhoff-von Neumann). Let X be a $p \times p$ bistochastic matrix. There exists an algorithm that runs in $O(p^{4.5})$ time and computes a decomposition $X = \sum_{k=1}^q \lambda^k \cdot A^k$, where $q \leq p^2 - p + 2$; for each $k \in [q]$, $\lambda^k \in [0, 1]$ and A^k is a $p \times p$ permutation matrix; and $\sum_{k=1}^q \lambda^k = 1$.

We are now in a position to present the PS-lottery algorithm (Algorithm 1). The high-level description of the algorithm is as follows. We first add some dummy goods (which every agent prefers less than any real good) to ensure that there are exactly nc goods. The expanded set of goods is called M' . Next, we simulate PS with this expanded set of goods M' . The algorithm runs for exactly c units of time because each agent eats exactly one unit of good per unit time. This produces a fractional allocation which is an $n \times (cn)$ matrix. To apply Theorem 1, we need to convert it into a square bistochastic matrix. For this, we track how much of each good each agent ate at each integral unit of time $[t-1, t]$, $t \in [c]$, while running PS. While agents eat one unit of good for each unit of time, note that this unit may consist of smaller fractions of several different goods. Then, we create a new set of agents $N' = \{i_1, \dots, i_c : i \in N\}$, where agents i_1, \dots, i_c represent agent i . We allocate the one unit of good eaten by agent i during time step $[t-1, t]$, $t \in [c]$, to the representative agent i_t . This produces a fractional allocation Y given by an $(nc) \times (nc)$ bistochastic matrix. We invoke Theorem 1 to decompose it into a convex combination of permutation matrices (in which each representative agent receives a single good). The permutation matrices are then modified by removing the dummy goods and combining the allocations to all representative agents $\{i_1, \dots, i_c\}$ back into an allocation to the agent i that they represent, for each $i \in N$. The convex combination over the modified permutation matrices gives us the desired solution, which is both ex ante EF and ex post EF1.

Algorithm 1 (PS-Lottery Algorithm)

- Input:** $I = (N, M, \succsim)$ where $|N| = n$, $|M| = m$.
Output: EF fractional allocation $X = \sum_{k=1}^K \lambda^k A^k$ where each A^k represents a deterministic EF1 allocation and $K \leq (m+n)^2 - 2(m+n) + 2$.
 1: $c \leftarrow \lceil m/n \rceil$.
 2: If m is a multiple of n , $D = \emptyset$. Else, $D = \{d_1, \dots, d_{nc-m}\}$.
 3: $M' \leftarrow M \cup D$ so that $|M'| = cn$.
 4: Set the preference profile \succsim' of agents in N over goods in M' as follows: for all $o, o' \in M$ and for all $i \in N$, $o \succsim'_i o'$ if $o \succsim_i o'$. For all $o \in M$ and $d \in D$, $o \succ'_i d$. The preferences between pairs of dummy goods can be arbitrary. {//see the forthcoming paragraph on “Additionally Achieving Efficiency” for how tie-breaking can impact efficiency}
 5: Run PS on the instance (N, M', \succsim') to get a fractional allocation X' .
 6: $N' \leftarrow \{i_1, \dots, i_c : i \in N\}$. Agents i_1, \dots, i_c are termed representatives of agent i .
 7: Construct a fractional allocation Y of goods in M' to agents in N' , where, for each $i \in N$ and $t \in [c]$, the goods eaten by agent i during time interval $[t-1, t]$ are allocated to its representative agent i_t . Note that Y is a $(cn) \times (cn)$ bistochastic matrix.
 8: Invoke Theorem 1 to compute a decomposition $Y = \sum_{k=1}^K \lambda^k B^k$ where $K \leq (cn)^2 - 2cn + 2$.
 9: Convert $Y = \sum_{k=1}^K \lambda^k B^k$ into $X = \sum_{k=1}^K \lambda^k A^k$ where all the dummy goods are ignored and each agent gets the allocation of its representatives.
 10: **return** Allocation X for instance I and its decomposition $\sum_{k=1}^K \lambda^k A^k$.

Lemma 1. Fix an integral allocation B^k from Line 8 of Algorithm 1 and let $t, t' \in [c]$. Let g_i^t denote the good allocated to representative agent i_t in B^k . Then agent i (weakly) prefers g_i^t to $g_{i'}^{t'}$ for any $i' \in N$ and $t' > t$ (that is, agent i weakly prefers the good allocated to its representative for time t to the good allocated to any representative of any agent for any time later than t).

Proof. Suppose, for contradiction, that there exists an $i' \in N$ and $t' > t$ such that $g_{i'}^{t'} \succ_i g_i^t$. By Theorem 1, representative agent i_t must have eaten a nonzero share of good g_i^t . That is, agent i must have eaten a nonzero share of good g_i^t in time period $[t-1, t]$. Similarly, representative agent $i'_{t'}$ must have eaten a nonzero share of good $g_{i'}^{t'}$ (in particular, agent i' eats a nonzero share of good $g_{i'}^{t'}$ in time period $[t'-1, t']$), and thus $g_{i'}^{t'}$ was not fully consumed at time t . However, because $g_{i'}^{t'} \succ_i g_i^t$, agent i should have fully consumed $g_{i'}^{t'}$ before consuming g_i^t , a contradiction. \square

Lemma 2. Every integral allocation returned by Algorithm 1 is envy-free up to one good.

Proof. Fix an integral allocation A^k from the output of Algorithm 1. For any $i \in [n]$ and $t \in [c]$, let g_i^t denote the good received by representative agent i_t in A^k (if

such a good exists). Fix a pair of agents $i, j \in N$. By Lemma 1, for every $t \in [c]$, we have $v_i(g_i^t) \geq v_i(g_j^{t+1})$. Hence,

$$v_i(A_i) \geq \sum_{t=1}^{c-1} v_i(g_i^t) \geq \sum_{t=2}^c v_i(g_j^t) = v_i(A_j \setminus \{g_j^1\}),$$

which establishes EF1. \square

Theorem 2. The randomized allocation implemented by Algorithm 1 is ex ante envy-free and ex post envy-free up to one good.

Proof. The ex post EF1 guarantee follows readily from Lemma 2. Ex ante EF follows from the fact that the fractional allocation X returned by the algorithm is equivalent to the fractional allocation returned by PS, thus inheriting its ex ante properties, including envy-freeness. The equivalence can be seen by noting that X is computed by taking the PS allocation X' (that includes dummy goods), distributing the goods allocated to agent i among its representative agents, and then recombining those goods back into a single bundle $X_i = X'_i \setminus D$. \square

Remark 1. The PS-lottery algorithm only uses the ordinal information. Therefore, its envy-freeness guarantees imply the same properties for all cardinal utilities consistent with the ordinal preferences. Hence, it returns an ex ante SD-envy-free and ex post SD-EF1 outcome.

Remark 2. The running time of Algorithm 1 is $O((m+n)^{4.5})$. The running time is dominated by the invocation of Theorem 1 on a $p \times p$ matrix, where $p = cn \leq m+n$, which takes $O(p^{4.5}) = O((m+n)^{4.5})$ time. The other operations such as computing an outcome of PS and transforming matrices take at most $O(nm)$ time.

Remark 3. Algorithm 1 is a combinatorial algorithm that computes a lottery over at most $(m+n)^2$ deterministic allocations. By Carathéodory’s theorem, any $n \times m$ fractional allocation that is represented by a convex combination of some K deterministic allocations can be represented by a convex combination of at most $nm + 1$ deterministic allocations among those K deterministic allocations. We can reduce the support of the lottery returned by Algorithm 1 to one involving at most $nm + 1$ deterministic SD-EF1 and SD-efficient allocations as follows. By using Gaussian elimination, we compute the subset of the set of matrices $\{A^1, \dots, A^K\}$ that forms the basis of A^1, \dots, A^K . We can then compute a convex combination of the matrices in the basis to achieve the same fractional allocation X .

We note that, whereas the PS-Lottery algorithm provides a way to implement PS by EF1 allocations, not every implementation of the PS outcome may satisfy ex post EF1. For example, consider the case of two agents with identical preferences over two goods. In that case, tossing a coin and then giving both goods to one agent is ex ante equivalent to the PS outcome. However, it is

not EF1 if agents have strictly positive utilities for both goods.

Next, we present a simple example showing how our algorithm works.

Example 2. Consider the instance in Example 1 in which $N = \{1, 2\}$, $M = \{g_1, g_2, g_3, g_4\}$. Suppose the valuations are as in Table 1.

The ordinal preferences of the agents over the goods are as follows: $v_1(g_1) > v_1(g_2) > v_1(g_3) > v_1(g_4)$ and $v_2(g_1) > v_2(g_3) > v_2(g_2) > v_2(g_4)$.

If we run the PS algorithm, we get the following outcome:

$$X' = \begin{bmatrix} 1/2 & 1 & 0 & 1/2 \\ 1/2 & 0 & 1 & 1/2 \end{bmatrix}.$$

It can be checked that X' is SD-envy-free. Because m is a multiple of n , $D = \emptyset$ and hence $M' = M \cup D = M$. We now show how to achieve our desired lottery to achieve the PS outcome. We run the PS rule on (N, M', \succ') to get allocation X . Then, for each agent's bundle, we let successive representative agents eat exactly one unit of goods one by one to get the following allocation, where i_j denotes the j th representative agent of i :

$$Y = \begin{matrix} 1_1 \\ 2_1 \\ 1_2 \\ 2_2 \end{matrix} \begin{bmatrix} 1/2 & 1/2 & 0 & 0 \\ 1/2 & 0 & 1/2 & 0 \\ 0 & 1/2 & 0 & 1/2 \\ 0 & 0 & 1/2 & 1/2 \end{bmatrix}$$

$$\stackrel{\text{BvN}}{=} \frac{1}{2} \cdot \begin{matrix} 1_1 \\ 2_1 \\ 1_2 \\ 2_2 \end{matrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} + \frac{1}{2} \cdot \begin{matrix} 1_1 \\ 2_1 \\ 1_2 \\ 2_2 \end{matrix} \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}.$$

Translating these for our original instance, we get the following decomposition over EF1 allocations:

$$X = \frac{1}{2} \cdot \begin{bmatrix} 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 \end{bmatrix} + \frac{1}{2} \cdot \begin{bmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{bmatrix}.$$

3.1. Additionally Achieving Efficiency

Because the outcome returned by Algorithm 1 is a lottery implementation of the PS rule, our algorithm inherits all the ex ante properties that the PS rule satisfies. When there are ties in the agents' preferences, it is

known that PS may fail to satisfy ex ante SD-Efficiency unless the ties are handled carefully. If we care about ex ante SD-Efficiency, then we do not artificially break any ties and can run the extended probabilistic serial (EPS) algorithm (Katta and Sethuraman 2006). The EPS algorithm uses the same continuous eating process as PS but, rather than breaking ties arbitrarily, makes coordinated choices for agents to eat one of their most preferred goods, using parametric network flows to compute the outcome. By doing so, it guarantees an outcome that is ex ante SD efficient and ex ante envy-free, even in the presence of ties. For $m \geq n$ goods, the algorithm takes time $O(m^3 \log m)$.

The formal specification of our *EPS-Lottery* algorithm is provided as Algorithm A.1 in Appendix A. It differs from Algorithm 1 only in Step 5, where it runs the EPS algorithm rather than the PS algorithm. The returned fractional allocation X is equivalent to the fractional allocation output by EPS and therefore inherits SD-Efficiency and envy-freeness. Additionally, the outcome can be implemented by EF1 deterministic allocations in the same manner as for the PS-lottery algorithm. To see this, the proof of Lemma 1 relies only on the fact that at every point in time, each agent eats (one of) their most preferred remaining goods. This property continues to hold when we use EPS as our base algorithm instead of PS. Finally, the running time is unchanged as well because the bottleneck step is to invoke Theorem 1, which takes time $O((m+n)^{4.5})$.

Theorem 3. *There is an algorithm that runs in time $O((m+n)^{4.5})$ and computes a randomized allocation that is ex ante SD-envy-free, ex ante SD efficient, and ex post SD-EF1.*

4. Impossibilities

In the previous section, we showed that ex ante (SD-)EF, ex post (SD-)EF1, and ex ante SD-Efficiency can be achieved simultaneously, thus providing a compelling solution for achieving both ex ante and ex post fairness in resource allocation problems. However, the SD-Efficiency guarantee is rather weak for settings where cardinal valuation functions are available. The obvious question then is whether we can achieve stronger efficiency guarantees along with ex ante and ex post fairness. Let us now consider the three cardinal efficiency notions from Section 2 that are related through the following logical implications: ex ante fPO \Rightarrow ex post fPO \Rightarrow ex post PO.

We first consider adding the weakest of them: ex post PO. Unfortunately, we were not able to settle whether ex ante EF (or even the weaker ex ante Prop) is compatible with ex post EF1 and ex post PO. This is the most compelling open question raised by our work, and we will return to it shortly. We can, however, show that strengthening ex ante EF to ex ante SD-EF immediately yields an incompatibility with ex post EF1 and ex post PO. Therefore, any algorithm

Table 1.

		Goods			
		g_1	g_2	g_3	g_4
Agents	1:	60	25	10	5
	2:	90	3	5	2

that is ex ante envy-free and based only on the ordinal preferences of the agents (such as those based on probabilistic serial) will necessarily fail either ex post EF1 or ex post PO.

Theorem 4. *There exists an instance with two agents and additive valuations in which no randomized allocation is simultaneously ex ante SD-envy-free, ex post envy-free up to one good, and ex post Pareto optimal.*

Proof. Consider the example in which $N = \{1, 2\}$, $M = \{a, b_1, b_2, b_3\}$ and the agents have utilities as given in Table 2.

The three goods b_1, b_2, b_3 are identical goods that we refer to as b goods. Ex ante SD-EF implies that each agent in expectation gets $1/2$ of a and 1.5 units of type b goods. Our first claim is that in any lottery implementing such an ex ante SD-EF allocation, there is at least one ex post allocation in which agent 2 must get good a . This follows from the fact that agent 2 gets $1/2$ of a in expectation.

Our second claim is that in any integral EF1 and PO allocation, agent 2 cannot get good a . Suppose for contradiction that agent 2 gets a . Then, EF1 requires that agent 1 gets at least two goods of type b . However, then, agent 1 can exchange these two goods for a to obtain a Pareto improvement.

From the previous two claims, it follows that for this instance, there exists no lottery over integral EF1 and PO outcomes that implements the SD-EF random outcome. \square

Let us move on and consider imposing a slightly stronger efficiency notion: ex post fPO. Integral EF1 + fPO allocations are known to always exist (Barman et al. 2018). Hence, the question of whether we can randomize over such allocations to achieve a desirable ex ante fairness guarantee is meaningful. However, in this case, we show that achieving ex ante EF is impossible along with ex post EF1 and ex post fPO.

Theorem 5. *There exists an instance with two agents and additive valuations in which no randomized allocation is simultaneously ex ante envy-free, ex post envy-free up to one good, and ex post fractionally Pareto optimal.*

Proof. We present an instance in which the unique integral allocation satisfying EF1 + fPO violates EF. Specifically, consider an instance with two goods (g_1, g_2) and two agents (1, 2) whose additive valuations are given in Table 3.

Table 2.

		Goods			
		a	b_1	b_2	b_3
Agents	1:	7	1	1	1
	2:	4	2	2	2

This instance has exactly two integral EF1 allocations: $A := (\{g_1\}, \{g_2\})$ and $B := (\{g_2\}, \{g_1\})$. It is easy to check that B is not fPO because it is Pareto dominated by a fractional allocation X that assigns g_1 completely to agent 1 and splits g_2 equally between the two agents. Indeed, $v_1(X_1) = v_1(g_1) + 0.5 \cdot v_1(g_2) = 2 \geq v_1(B_1)$ and $v_2(X_2) = 0.5 \cdot v_2(g_2) = 1.5 > v_2(B_2)$. To see why A is fPO, notice that it assigns each good to an agent that has the highest valuation for it. Therefore, A maximizes the utilitarian social welfare (i.e., sum of agents' utilities), which implies that it is fPO. Therefore, any randomized allocation that is ex post EF1 + fPO must be supported entirely on the integral allocation A . However, A violates EF because $v_1(A_1) = v_1(g_1) < v_1(g_2) = v_1(A_2)$. Therefore, the randomized allocation is not ex ante envy-free. \square

In the proof of Theorem 5, not only does allocation A violate envy-freeness, but it also violates the weaker property of proportionality. Therefore, Theorem 5 continues to hold even when we replace ex ante EF by ex ante Prop.

We mentioned earlier that achieving ex ante EF along with ex post EF1 + PO is an open question. Two prominent methods for finding an integral EF1+PO allocation are the integral MNW rule (Caragiannis et al. 2019), which maximizes the product of agents' utilities, and the market-based rule of Barman et al. (2018). An interesting implication of Theorem 5 is that we cannot hope to achieve ex ante envy-freeness (or even ex ante proportionality) by randomizing over allocations returned by either method. The latter method is guaranteed to return an integral EF1 + fPO allocation, so Theorem 5 directly applies. The MNW rule, although only guaranteed to return an integral EF1 + PO allocation, uniquely returns allocation A in the example presented in the proof of Theorem 5, which violates proportionality (and therefore envy-freeness).

Finally, when we consider the strongest efficiency property of ex ante fPO, we find that it is incompatible with ex ante SD-EF even without imposing an ex post fairness guarantee. The theorem follows directly from theorem 5 of Aziz and Ye (2014), but we reprove it in our context for the sake of completeness.

Theorem 6. *There exists an instance with two agents and additive valuations in which no randomized allocation is simultaneously ex ante SD-envy-free and ex ante fractionally Pareto optimal.*

Table 3.

		Goods	
		g_1	g_2
Agents	1:	1	2
	2:	1	3

Proof. Consider the same instance as in the proof of Theorem 5, with two goods (g_1, g_2) and two agents $(1, 2)$ whose additive valuations are $v_1(g_1) = 1, v_1(g_2) = 2$ and $v_2(g_1) = 1, v_2(g_2) = 3$. Because both agents prefer g_2 to g_1 , the only allocation X satisfying SD-EF is the one that allocates each agent half of each good, with $v_1(X) = 1.5$ and $v_2(X) = 2$. However, this allocation is Pareto dominated by the allocation Y that allocates g_1 and one third of g_2 to agent 1 and allocates two thirds of g_2 to agent 2, with $v_1(Y) = 5/3 > v_1(X)$ and $v_2(Y) = 2 = v_2(X)$. \square

5. MNW-Lottery Algorithm

Given the impossibilities in the previous section, together with the difficulty of achieving ex ante EF and ex post EF1 + PO using known techniques, it is evident that achieving any efficiency property stronger than SD-Efficiency requires relaxing at least one of the fairness guarantees. In this section, we focus on relaxing ex post EF1. There are two prominent relaxations that have been proposed in the literature—namely, Prop1 and EF₁¹.

Proportionality up to one good requires that an agent be able to achieve their proportional share by adding a single good to their bundle.

Definition 8 (Proportionality up to One Good (Prop1); Conitzer et al. 2017). An integral allocation A is *proportional up to one good* if for every agent $i \in N$, either $v_i(A_i) \geq v_i(M)/n$ or there exists a good $g \notin A_i$ such that $v_i(A_i \cup \{g\}) \geq v_i(M)/n$.

The next property, called envy-freeness up to one good more-and-less (EF₁¹) (Barman and Krishnamurthy 2019), is a relaxation of EF1 and enjoys strong algorithmic support in conjunction with PO (Barman and Krishnamurthy 2019). It allows a good to be removed from the envied agent's bundle, and a (possibly different) good to be added to the envying agent's.

Definition 9 (Envy-Freeness up to One Good More-and-Less (EF₁¹)). An integral allocation A is *envy-free up to one good more-and-less* if for every pair of agents $i, j \in N$ such that $A_j \neq \emptyset$, we have $v_i(A_i \cup \{g_i\}) \geq v_i(A_j \setminus \{g_j\})$ for some goods $g_i \notin A_i$ and $g_j \in A_j$.

Prop1 and EF₁¹ are incomparable in general, in the sense that allocations satisfying one property may not satisfy the other.

We show that both can be achieved simultaneously, and in fact, this can be done while also achieving one of the strongest ex ante properties called ex ante group fairness. Group fairness simultaneously strengthens various properties including envy-freeness and PO by offering fairness guarantees to groups of arbitrary size and composition. Given any set $S \subseteq N$ of agents, we write $\cup_{i \in S} X_i$ to denote the union of the fractional allocations to agents in S , that is, $\cup_{i \in S} X_i := (\sum_{i \in S} X_{i,g})_{g \in M}$.

Definition 10 (Group Fairness (GF); Conitzer et al. 2019). An allocation X is *group fair* if for all nonempty subsets of agents $S, T \subseteq N$, there is no fractional allocation Y of $\cup_{i \in T} X_i$ to the agents in S such that $\frac{|S|}{|T|} \cdot v_i(Y_i) \geq v_i(X_i)$ for all agents $i \in S$ and at least one inequality is strict.

Imposing the previous constraint over restricted (S, T) pairs can recover properties such as proportionality ($|S| = 1, T = N$), envy-freeness ($|S| = |T| = 1$), and fractional PO ($S = T = N$). Because group fairness implies envy-freeness, it is clear that it may not be possible to achieve ex post group fairness, so we will only focus on ex ante group fairness. Ex ante GF not only implies ex ante envy-freeness, but also ex ante fPO, which, by Proposition 1, implies ex post fPO for any implementation of it. In other words, our goal is to implement a fractional GF allocation using integral Prop1 + EF₁¹ allocations.

Before we explain how we achieve the properties defined previously, we define an important fractional allocation rule that will be central to the study in this section.

Definition 11 (Fractional MNW Rule). Given an instance $I \in \mathcal{I}$, the fractional MNW rule returns all fractional allocations that maximize the product of agents' utilities, that is, $\text{MNW}(I) := \arg \max_{X \in \mathcal{X}} \prod_{i \in N} v_i(X_i)$. We refer to an allocation $A \in \text{MNW}(I)$ as a fractional MNW allocation.

It is known that any fractional MNW allocation satisfies group fairness (Conitzer et al. 2019). Furthermore, we know that a fractional MNW allocation can be computed in strongly polynomial time (Orlin 2010, Végh 2016). Hence, we ask whether a fractional MNW allocation can be implemented using integral Prop1 + EF₁¹ allocations. Our starting point is a result by Budish et al. (2013) that allows implementing any fractional allocation using integral allocations that are very "close" to it in agent utilities. Specifically, they prove the next result deriving and using an extension of the Birkhoff-von Neumann theorem (Theorem 1).

Proposition 2 (Utility Guarantee; Theorem 9 of Budish et al. 2013). *Given any fractional allocation X , one can compute, in strongly polynomial time, a randomized allocation implementing X whose support consists of integral allocations A^1, \dots, A^K such that for every $k \in [K]$ and every agent $i \in N$,*

$$|v_i(X_i) - v_i(A_i^k)| \leq \max\{v_i(g) - v_i(g') : 0 < X_{i,g}, X_{i,g'} < 1\}.$$

The upper bound established in Proposition 2 on how much agent i 's utility under an integral allocation A^k in the support can differ from their utility under the fractional allocation X depends only on their own fractional allocation X_i . In contrast, the fairness guarantees we want to establish for the integral allocations in the support—Prop1 and EF₁¹—consider what happens when we add a good to the bundle of agent i that agent i is not already allocated in the integral allocation A^k ; in other

words, we need a stronger guarantee for integral allocations in the support that depends on which goods the agent is (or is not) allocated ex post.

It turns out that the method proposed by Budish et al. (2013) already provides such a guarantee, and their proof can be adapted to establish a more nuanced bound. Specifically, we show that if the agent's ex ante utility $v_i(X_i)$ exceeds their ex post utility $v_i(A_i^k)$, then the gap is at most the maximum value the agent has for any good that they *lost* in the integral allocation (i.e., any good g such that $0 < X_{i,g} < 1$ and $A_{i,g}^k = 0$). Similarly, if the ex post utility exceeds the ex ante utility, then the gap is at most the maximum value the agent has for any good that they *gained* in the integral allocation (i.e., any good g such that $0 < X_{i,g} < 1$ and $A_{i,g}^k = 1$). We later show that this subtle improvement helps us establish the desired ex post fairness guarantees.

Lemma 3 (Utility Guarantee++). *Given a fractional allocation X , one can compute, in strongly polynomial time, a randomized allocation implementing X whose support consists of integral allocations A^1, \dots, A^K such that for every $k \in [K]$ and every agent $i \in N$, the following hold:*

1. If $v_i(A_i^k) < v_i(X_i)$, then $\exists g_i^- \notin A_i^k$ with $X_{i,g_i^-} > 0$ such that $v_i(A_i^k) + v_i(g_i^-) > v_i(X_i)$.
2. If $v_i(A_i^k) > v_i(X_i)$, then $\exists g_i^+ \in A_i^k$ with $X_{i,g_i^+} < 1$ such that $v_i(A_i^k) - v_i(g_i^+) < v_i(X_i)$.

The proof of Lemma 3 is presented in Section B.2 of Appendix B.

We now show how Lemma 3 can be used to achieve our desired ex post fairness guarantees of Prop1 and EF_1^1 . Our overall approach is summarized as Algorithm 2 and is referred to as the MNW-lottery algorithm.

Algorithm 2 (MNW-Lottery Algorithm)

Input: $I = (N, M, v)$.

Output: Fractional allocation $X = \sum_{k=1}^K \lambda^k A^k$ where each A^k represents an integral allocation.

- 1: $X \leftarrow$ Fractional MNW allocation (using an algorithm of Orlin (2010) or Végé (2016)).
- 2: For any $i \in N$ and any $k \in [m]$, let $Q_{i,k} := \sum_{t=1}^k X_{i,g_{i,t}}$ be the total fractional amount of the k most preferred goods assigned to agent i under X .
- 3: Consider the following set of bihierarchical constraints on a generic fractional allocation Y :

$$\begin{aligned} \mathcal{H}_1 : [Q_{i,k}] &\leq \sum_{t=1}^k Y_{i,g_{i,t}} \leq [Q_{i,k}], \quad \forall i \in N \text{ and } \forall k \in [m], \\ \mathcal{H}_2 : \sum_{i \in N} Y_{i,g} &= 1, \quad \forall g \in M. \end{aligned} \tag{1}$$

- 4: Use the Algorithm of Budish et al. (2013) (specified in their appendix B) to find the randomized allocation $\sum_{k=1}^K \lambda^k A^k$ implementing the fractional allocation X that satisfies the same constraints as (1).
- 5: **return** Allocation X for instance I and its decomposition $\sum_{k=1}^K \lambda^k A^k$.

Theorem 7. *There is a strongly polynomial-time algorithm that, given any fractional proportional allocation as input, computes an implementation of it using integral allocations that are proportional up to one good. If, in addition, the input is a fractional MNW allocation, then the integral allocations in the support also satisfy envy-freeness up to one good more-and-less.*

Proof. Let X be a fractional allocation and let A^1, \dots, A^K be integral allocations in the support of an implementation of X produced by Lemma 3.

Suppose X satisfies proportionality. We want to show that for each $k \in [K]$, A^k is Prop1. Because X is proportional, for every $i \in N$, $v_i(X_i) \geq v_i(M)/n$. Fix $k \in [K]$. By Lemma 3, we have that for every agent $i \in N$, either $v_i(A_i^k) \geq v_i(X_i) \geq v_i(M)/n$, or there exists a good $g \notin A_i^k$ such that $v_i(A_i^k) + v_i(g) > v_i(X_i) \geq v_i(M)/n$. Therefore, A^k is Prop1.

Next, suppose that X maximizes the Nash social welfare among all fractional allocations. Because a fractional MNW allocation is certainly proportional (Varian 1974), the aforementioned argument still applies for ex post Prop1. We show that in this case, A^k is also EF_1^1 for each $k \in [K]$. Because X is a fractional MNW allocation, the following condition is satisfied for any pair of agents $i, j \in N$ and any good $g \in M$ (the condition that transferring an arbitrarily small fraction of good g from agent i to agent j does not increase Nash welfare reduces to this condition):

$$X_{i,g} > 0 \implies \frac{v_i(g)}{v_i(X_i)} \geq \frac{v_j(g)}{v_j(X_j)}. \tag{2}$$

Fix a pair of distinct agents $i, j \in N$. By Lemma 3, either $v_i(A_i^k) \leq v_i(X_i)$, or there exists $g_i^+ \in A_i^k$ with $X_{i,g_i^+} < 1$ such that $v_i(A_i^k \setminus \{g_i^+\}) < v_i(X_i)$. Similarly, either $v_j(A_j^k) \geq v_j(X_j)$, or there exists $g_j^- \notin A_j^k$ with $X_{j,g_j^-} > 0$ such that $v_j(A_j^k \cup \{g_j^-\}) > v_j(X_j)$. To simplify the analysis, let us assume that the second condition holds in both cases. (If $v_i(A_i^k) \leq v_i(X_i)$ (or $v_j(A_j^k) \geq v_j(X_j)$), we can treat g_i^+ (or g_j^-) as a dummy good with $v_i(g_i^+) = 0$ (or $v_j(g_j^-) = 0$)).

By summing the right-hand side inequality in Equation (2) over all $g \in A_i^k \setminus \{g_i^+\}$, we get

$$\frac{v_j(A_i^k \setminus \{g_i^+\})}{v_j(X_j)} \leq \frac{v_i(A_i^k \setminus \{g_i^+\})}{v_i(X_i)} < 1.$$

Thus, $v_j(A_i^k \setminus \{g_i^+\}) < v_j(X_j) < v_j(A_j^k \cup \{g_j^-\})$, implying that A^k satisfies EF_1^1 , as desired. \square

Remark 4. The proof of Theorem 7 establishes a stronger version of Prop1 wherein an agent not receiving their proportional share gets *strictly* more than their proportional share by receiving one additional good. Similarly, it also establishes a stronger version of EF_1^1 wherein an agent envying another agent would strictly prefer their own allocation over the other agent's

allocation after adding one missing good to their bundle and removing one good from the other agent's bundle.

Barman and Krishnamurthy (2019) recently established that integral Prop1 + EF₁¹ + fPO allocations exist and can be computed in strongly polynomial time. They rely on special-purpose techniques for rounding a fractional MNW allocation. By contrast, Theorem 7 uses a standard technique to round a fractional MNW allocation, computes not just *one* integral Prop1 + EF₁¹ + fPO allocation but rather an implementation of the fractional MNW allocation over such integral allocations, and can be applied to *any* fractional Prop + PO allocation to implement it using integral Prop1 + fPO allocations. Recall that because a fractional MNW allocation is Pareto optimal, any allocation in the support of an implementation of it must be fPO by Proposition 1.

In addition to guaranteeing EF₁¹ integral allocations in Theorem 7, the fractional MNW allocation is known to be group fair. This observation, along with the fact that a fractional MNW allocation can be computed in strongly polynomial time (Orlin 2010, Végh 2016), and Theorem 7 immediately yield the main result of this section.

Theorem 8. *There exists a randomized allocation that is ex ante group fair, ex post proportional up to one good, and ex post envy-free up to one good more-and-less. Furthermore, it can be computed in strongly polynomial time.*

Example 3. Consider again the instance in Example 1 in which $N = \{1, 2\}$, $M = \{g_1, g_2, g_3, g_4\}$, and the valuations of the agents for the goods are given in Table 4.

The unique fractional MNW allocation is

$$X = \begin{bmatrix} 1/6 & 1 & 1 & 1 \\ 5/6 & 0 & 0 & 0 \end{bmatrix}.$$

The unique randomized implementation of fractional allocation X (and, in particular, the randomized allocation output by the algorithm of Budish et al. (2013)) is

$$X = \frac{1}{6} \cdot \begin{bmatrix} 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} + \frac{5}{6} \cdot \begin{bmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 \end{bmatrix}.$$

Both integral allocations in the decomposition of X satisfy Prop1 and EF₁¹.

Table 4.

		Goods			
		g_1	g_2	g_3	g_4
Agents	1:	60	25	10	5
	2:	90	3	5	2

6. Discussion

We consider the question of how well fairness and efficiency can be satisfied from ex post and ex ante perspectives simultaneously. Although we focused on the case of allocating goods to agents, an interesting extension is the case where we instead have a set of negatively valued *bads* to allocate (e.g., chores to family members or committee assignments to department faculty). With appropriately generalized definitions, all our results continue to hold for the bads setting.

Algorithms 1 and 2 output a distribution over integral allocations rather than simply producing a sample. Although a sample may be sufficient for making an allocation, it has a key limitation: it is impossible for a participant to audit the ex ante properties of the distribution. On the other hand, from a transparency perspective, publishing a distribution allows ex ante properties to be verified, as long as participants trust the mechanism by which a random sample is chosen.

Perhaps the most fascinating open question that stems from our work is whether ex ante envy-freeness (or even ex ante proportionality) is compatible with ex post EF1 and ex post PO.

Open Question: Does there always exist a randomized allocation that is ex ante EF, ex post EF1, and ex post PO? What about ex ante Prop, ex post EF1, and ex post PO?

The difficulty in approaching this question is that there are very few available methods of finding integral EF1 + PO allocations (Barman et al. 2018, Caragiannis et al. 2019), so finding many such allocations and randomizing over them is tricky. Also, unlike the set of integral EF1 allocations, which we somewhat understand, not much is known about the set of integral EF1 + PO allocations other than the fact that it is always nonempty.

Various other open problems remain. For instance, other fairness concepts such as envy-freeness up to any good (EFX) (Caragiannis et al. 2019) or approximate maximin share fairness (MMS) (Budish 2011) can be considered. Future work can also consider the price of fairness: What fraction of the optimal social welfare must be sacrificed (in the worst case) to guarantee a fair allocation? Bertsimas et al. (2011) and Caragiannis et al. (2012) study the price of EF, whereas Bei et al. (2021) and Barman et al. (2020) study the price of EF1 for indivisible goods. What is the price of achieving both ex ante EF and ex post EF1 together?

More broadly, the next step would be to achieve ex ante and ex post fairness guarantees simultaneously in a variety of other problems such as voting, matching, and public decision making.

Acknowledgments

The authors thank the anonymous reviewers from EC'20, Jeremy Lindsay, and Dominik Peters for helpful comments.

This paper is an extended version that combines the main results of earlier work by Freeman et al. (2020) and Aziz (2020).

Appendix A. EPS-Lottery Algorithm

Algorithm A.1 (EPS-Lottery Algorithm)

- Input:** $I = (N, M, \succ)$ where $|N| = n$, $|M| = m$.
Output: EF fractional allocation $X = \sum_{k=1}^K \lambda^k A^k$ where each A^k represents a deterministic EF1 allocation and $K \leq (m+n)^2 - 2(m+n) + 2$.
 1: $c \leftarrow \lceil m/n \rceil$.
 2: If m is a multiple of n , $D = \emptyset$. Else, $D = \{d_1, \dots, d_{nc-m}\}$.
 3: $M' \leftarrow M \cup D$ so that $|M'| = cn$.
 4: Set the preference profile \succ' of agents in N over goods in M' as follows: for all $o, o' \in M$ and for all $i \in N$, $o \succ'_i o'$ if $o \succ_i o'$; for all $o \in M$ and $d \in D$, $o \succ'_i d$. The preferences between pairs of dummy goods can be arbitrary.
 5: Run EPS on the instance (N, M', \succ') to get a fractional allocation X' .
 6: $N' \leftarrow \{i_1, \dots, i_c : i \in N\}$. Agents i_1, \dots, i_c are termed as representatives of agent i .
 7: Construct a fractional allocation Y of goods in M' to agents in N' , where, for each $i \in N$ and $t \in [c]$, the goods eaten by agent i during time interval $[t-1, t]$ are allocated to its representative agent i_t . Note that Y is a $(cn) \times (cn)$ bistochastic matrix.
 8: Invoke Theorem 1 to compute a decomposition $Y = \sum_{k=1}^K \lambda^k B^k$, where $K \leq (cn)^2 - 2cn + 2$.
 9: Convert $Y = \sum_{k=1}^K \lambda^k B^k$ into $X = \sum_{k=1}^K \lambda^k A^k$ where all the dummy goods are ignored and each agent gets the allocation of its representatives.
 10: **return** Allocation X for instance I and its decomposition $\sum_{k=1}^K \lambda^k A^k$.

Appendix B. Omitted Material from Section 5

B.1. Decomposition Result of Budish et al. (2013)

Let X be a fractional allocation. Recall that X satisfies *column-wise* feasibility constraints, namely $0 \leq \sum_{i \in N} X_{i,r} \leq 1$ for all $r \in M$. More generally, we can impose capacity constraints of the form $q_S \leq \sum_{(i,r) \in S} X_{i,r} \leq \bar{q}_S$, where S is a *constraint set* comprising of a collection of agent-object pairs, and q_S and \bar{q}_S are the *lower* and *upper quotas* for S , respectively. The set of all capacity constraints imposed by a given problem is called the *constraint structure* \mathcal{H} of the problem and is specified as a collection of all constraint sets and the corresponding quotas $(q_S, \bar{q}_S)_{S \in \mathcal{H}}$. Given a constraint structure \mathcal{H} , we say that the fractional allocation X admits a *feasible* implementation $\mathbf{X} := \{(p^k, A^k)\}_{k \in [\ell]}$ if every integral allocation in its

support also satisfies the constraints in \mathcal{H} . That is, for every $k \in [\ell]$, we have

$$q_S \leq \sum_{(i,r) \in S} A^k_{i,r} \leq \bar{q}_S \text{ for every } S \in \mathcal{H}.$$

Definition B.1 (Hierarchy and Bihierarchy). A constraint structure \mathcal{H} is said to be a *hierarchy* (or a laminar family) if for every $S, S' \in \mathcal{H}$, we have that either $S \subset S'$, or $S' \subset S$, or $S \cap S' = \emptyset$. We say that \mathcal{H} is a *bihierarchy* if it can be partitioned into two hierarchies, that is, if there exist hierarchies \mathcal{H}_1 and \mathcal{H}_2 such that $\mathcal{H} = \mathcal{H}_1 \cup \mathcal{H}_2$ and $\mathcal{H}_1 \cap \mathcal{H}_2 = \emptyset$.

As an example, consider the fractional allocation X in Figure B.1. The row constraints (shown as red or blue solid rectangles) and all singleton constraints of the form $0 \leq X_{i,r} \leq 1$ (not shown in the figure) together constitute a hierarchy, say \mathcal{H}_1 , because for any pair of constraint sets, either they are disjoint or one is completely contained inside the other. Similarly, the column constraints (shown as gray dotted rectangles) form another hierarchy \mathcal{H}_2 . Furthermore, $\mathcal{H} := \mathcal{H}_1 \cup \mathcal{H}_2$ is a bihierarchy because any constraint set (rectangle or singleton) belongs to exactly one of \mathcal{H}_1 or \mathcal{H}_2 .

Budish et al. (2013) showed that bihierarchy constraint structure is a sufficient condition for a fractional allocation to admit a feasible implementation. We recall this result in Proposition B.1.

Proposition B.1 (Budish et al. 2013). *Given a fractional allocation X satisfying a bihierarchy constraint structure \mathcal{H} , one can compute, in strongly polynomial time, a set of coefficients $p_1, \dots, p_\ell \in [0, 1]$ and integral allocations A^1, \dots, A^ℓ such that (a) $\sum_{k=1}^\ell p^k = 1$, (b) each A^k satisfies the constraints in \mathcal{H} , and (c) $X = \sum_{k=1}^\ell p^k A^k$.*

Observe that the well-known Birkhoff-von Neumann theorem is a special case of Proposition B.1 when \mathcal{H}_1 consists of all singleton and row constraints, \mathcal{H}_2 consists of all column constraints. The lower and upper quotas for the singleton constraints are zero and one, respectively, whereas those for the row and column constraints are one and one. It is worth pointing out that, although Budish et al. (2013) only note a polynomial running time, it is easy to check that their (combinatorial) algorithm, in fact, runs in *strongly* polynomial time.


Budish et al. (2013) use Proposition B.1 to establish the utility guarantee (Proposition 2).

B.2. Proof of Lemma 3

Recall the statement of Lemma 3.

Lemma 3 (Utility Guarantee++). *Given a fractional allocation X , one can compute, in strongly polynomial time, a randomized*

Figure B.1. Decomposition of a Fractional Proportional Allocation X into Integral Prop1 Allocations A^1, A^2 , and A^3



Note. The underlying fair division instance comprises four goods and two agents with valuations $v_1(1) = 10, v_1(2) = 6, v_1(3) = 4, v_1(4) = 2$ and $v_2(1) = 2, v_2(2) = 10, v_2(3) = 6, v_2(4) = 4$.

allocation implementing X whose support consists of integral allocations A^1, \dots, A^K such that for every $k \in [K]$ and every agent $i \in N$, the following hold:

1. If $v_i(A_i^k) < v_i(X_i)$, then $\exists g_i^- \notin A_i^k$ with $X_{i,g_i^-} > 0$ such that $v_i(A_i^k) + v_i(g_i^-) > v_i(X_i)$.
2. If $v_i(A_i^k) > v_i(X_i)$, then $\exists g_i^+ \in A_i^k$ with $X_{i,g_i^+} < 1$ such that $v_i(A_i^k) - v_i(g_i^+) < v_i(X_i)$.

Proof. In their proof of proposition 2, Budish et al. (2013) propose the following method for computing an implementation of a given fractional allocation X . Consider a fixed agent $i \in N$. Suppose the goods in M are indexed as $g_{i,1}, \dots, g_{i,m}$ so that $v_i(g_{i,k}) \geq v_i(g_{i,k+1})$ for each $k \in [m-1]$. For simplicity, we will write $v_i(k) := v_i(g_{i,k})$ for all $k \in [m]$ and $v_i(m+1) := 0$.

For any $i \in N$ and any $k \in [m]$, define $Q_{i,k} := \sum_{t=1}^k X_{i,g_{i,t}}$ as the total fractional amount of the k most preferred goods assigned to agent i under X . Consider the following set of constraints on a generic fractional allocation Y (for simplicity, we omit the singleton constraints defining a valid allocation, $0 \leq Y_{i,g} \leq 1$ for all $i \in N$ and $g \in M$, although they can be included in either hierarchy without violating the hierarchy structure):

$$\begin{aligned} \mathcal{H}_1 : [Q_{i,k}] &\leq \sum_{t=1}^k Y_{i,g_{i,t}} \leq \lceil Q_{i,k} \rceil, \quad \forall i \in N \text{ and } \forall k \in [m], \\ \mathcal{H}_2 : \sum_{i \in N} Y_{i,g} &= 1, \quad \forall g \in M. \end{aligned} \quad (\text{B.1})$$

For illustrative purposes, the solid red and blue rectangles in Figure B.1 correspond to \mathcal{H}_1 , and the dotted gray rectangles correspond to \mathcal{H}_2 . Observe that X trivially satisfies these constraints. Budish et al. show that these constraints have the so-called “bihierarchy” structure (refer to Section B.1 for a formal definition), which allows computing, in strongly polynomial time, an implementation of X whose support consists of integral allocations A^1, \dots, A^l that also satisfy these constraints.

Last, Budish et al. show that any integral allocation satisfying the constraints in Equation (B.1) must satisfy the guarantee in Proposition 2. We show that it in fact satisfies the slightly stronger guarantee that we seek. For simplicity, let us write \hat{A} to denote a generic integral allocation satisfying the constraints in Equation (B.1), and $\hat{Q}_{i,k} := \sum_{t=1}^k \hat{A}_{i,g_{i,t}}$ for all $i \in N$ and $k \in [m]$.

Let us first analyze the case where $v_i(\hat{A}_i) < v_i(X_i)$. Then, there must exist some good $g \in M$ such that $\hat{A}_{i,g} < X_{i,g}$. Because $\hat{A}_{i,g} \in \{0, 1\}$ and $X_{i,g} \in [0, 1]$, this is equivalent to $\hat{A}_{i,g} = 0 < X_{i,g}$. Let k^- be the smallest index such that $\hat{A}_{i,g_{i,k^-}} < X_{i,g_{i,k^-}}$; that is, g_{i,k^-} is agent i 's most preferred good satisfying this condition. Hence, $g_{i,k^-} \notin \hat{A}_i$ and $X_{i,g_{i,k^-}} > 0$. Furthermore, for all $k < k^-$, we have $\hat{A}_{i,g_{i,k}} \geq X_{i,g_{i,k}}$, and, as a result, $\hat{Q}_{i,k} \geq Q_{i,k}$. Thus,

$$\begin{aligned} v_i(X_i) - v_i(\hat{A}_i) &= \sum_{k=1}^m v_i(k) \cdot (X_{i,g_{i,k}} - \hat{A}_{i,g_{i,k}}) \\ &= \sum_{k=1}^m (v_i(k) - v_i(k+1)) \cdot (Q_{i,k} - \hat{Q}_{i,k}) \\ &\leq \sum_{k=k^-}^m (v_i(k) - v_i(k+1)) \cdot (Q_{i,k} - \hat{Q}_{i,k}) \\ &< \sum_{k=k^-}^m (v_i(k) - v_i(k+1)) \cdot 1 = v_i(k^-), \end{aligned}$$

where the second transition is a simple algebraic exercise, the third transition holds because we noted that $\hat{Q}_{i,k} \geq Q_{i,k}$ for all $k < k^-$, and the fourth transition holds because \hat{A} satisfies \mathcal{H}_1 in Equation (B.1), and therefore, we have that $Q_{i,k} - \hat{Q}_{i,k} \leq Q_{i,k} - \lfloor Q_{i,k} \rfloor < 1$. Taking $g_i^- := g_{i,k^-}$, this is the guarantee we desire when $v_i(\hat{A}_i) < v_i(X_i)$.

Next, consider the other case where $v_i(\hat{A}_i) > v_i(X_i)$. Then, there must exist some good $g \in M$ such that $\hat{A}_{i,g} > X_{i,g}$. This is equivalent to $\hat{A}_{i,g} = 1 > X_{i,g}$. Let k^+ be the smallest index such that g_{i,k^+} satisfies this condition. Then, we have that $g_{i,k^+} \in \hat{A}_i$, $X_{i,g_{i,k^+}} < 1$, and by an argument similar to the previous one, $v_i(\hat{A}_i) - v_i(X_i) < v_i(k^+)$. Hence, in this case, we can take $g_i^+ := g_{i,k^+}$, as desired. \square

References

- Abdulkadiroglu A, Sönmez T (1998) Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66(3):689–701.
- Akbarpour M, Nikzad A (2020) Approximate random allocation mechanisms. *Rev. Econom. Stud.* 87(6):2473–2510.
- Aleksandrov M, Aziz H, Gaspers S, Walsh T (2015) Online fair division: Analysing a food bank problem. Yang Q, Wooldridge M, eds. *Proc. 24th Internat. Joint Conf. on Artificial Intelligence (AAAI Press, Palo Alto, CA)*, 2540–2546.
- Alon N (1987) Splitting necklaces. *Adv. Math.* 63(3):247–253.
- Aziz H, Ye C (2014) Cake cutting algorithms for piecewise constant and piecewise uniform valuations. Liu TY, Qi Q, Ye Y, eds. *Proc. 10th Internat. Conf. on Web and Internet Econom.* (Springer, Berlin), 1–14.
- Aziz H (2019) A probabilistic approach to voting, allocation, matching, and coalition formation. *The Future of Economic Design* (Springer, Berlin), 45–50.
- Aziz H (2020) Simultaneously achieving ex-ante and ex-post fairness. Chen X, Gravin N, Hoefer M, Mehta R, eds. *Proc. 16th Internat. Conf. on Web and Internet Econom.* (Springer, Berlin), 341–355.
- Babaioff M, Ezra T, Feige U (2022) On best-of-both-worlds fair-share allocations. Hansen KA, Liu TX, Malekian A, eds. *18th Internat. Conf. Web Internet Econom.* (Springer, Berlin), 237–255.
- Barman S, Krishnamurthy S (2019) On the proximity of markets with integral equilibria. Hentenryck PV, Zhou Z-H, eds. *Proc. 33rd AAAI Conf. on Artificial Intelligence (AAAI Press, Palo Alto, CA)*, 1748–1755.
- Barman S, Bhaskar U, Shah N (2020) Optimal bounds on the price of fairness for indivisible goods. Chen X, Gravin N, Hoefer M, Mehta R, eds. *Proc. 16th Conf. on Web and Internet Econom.* (Springer, Berlin), 356–369.
- Barman S, Krishnamurthy SK, Vaish R (2018) Finding fair and efficient allocations. Elkind E, Vohra R, Tardos E, eds. *Proc. 19th Conf. on Econom. and Comput.* (ACM, New York), 557–574.
- Bei X, Lu X, Manurangsi P, Suksompong W (2021) The price of fairness for indivisible goods. *Theory Comput. Systems* 65: 1069–1093.
- Bertsimas D, Farias VF, Trichakis N (2011) The price of fairness. *Oper. Res.* 59(1):17–31.
- Birkhoff G (1946) Three observations on linear algebra. *Universidad Nacional Tucumán Rev. A* 5:147–151.
- Bogomolnaia A, Moulin H (2001) A new solution to the random assignment problem. *J. Econom. Theory* 100:295–328.
- Bouveret S, Chevaleyre Y, Maudet N (2016) Fair allocation of indivisible goods. Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD, eds. *Handbook of Computational Social Choice* (Cambridge University Press, Cambridge, United Kingdom), 284–310.
- Brams SJ, Taylor AD (1996) *Fair Division: From Cake-Cutting to Dispute Resolution* (Cambridge University Press, Cambridge, UK).

- Brandt F, Conitzer V, Endress U, Lang J, Procaccia AD, eds. (2016) *Handbook of Computational Social Choice* (Cambridge University Press, Cambridge, United Kingdom).
- Brustle J, Dippel J, Narayan V, Suzuki M, Vetta A (2020) One dollar each eliminates envy. Ostrovsky M, Procaccia A, Biro P, Hartline J, eds. *Proc. 21st ACM Conf. on Econom. and Comput* (ACM, New York), 23–39.
- Budish E (2011) The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *J. Political Econom.* 119(6):1061–1103.
- Budish E, Che Y-K, Kojima F, Milgrom P (2013) Designing random allocation mechanisms: Theory and applications. *Amer. Econom. Rev.* 103(2):585–623.
- Caragiannis I, Kanellopoulos P, Kyropoulou M (2021) On interim envy-free allocation lotteries. Chawla S, Echenique F, Biro P, Sodomka E, eds. *Proc. 22nd ACM Conf. on Econom. and Comput.* (ACM, New York), 264–284.
- Caragiannis I, Kaklamanis C, Kanellopoulos P, Kyropoulou M (2012) The efficiency of fair division. *Theory Comput. Systems* 50(4):589–610.
- Caragiannis I, Kurokawa D, Moulin H, Procaccia A, Shah N, Wang J (2019) The unreasonable fairness of maximum Nash welfare. *ACM Trans. Econom. Comput.* 7(3):1–32.
- Cechlárová K, Dobos J, Pillárová E (2013) On the existence of equitable cake divisions. *Inform. Sci.* 228:239–245.
- Chen Y, Sönmez T (2002) Improving efficiency of on-campus housing: An experimental study. *Amer. Econom. Rev.* 92(5):1669–1686.
- Conitzer V, Freeman R, Shah N (2017) Fair public decision making. Babaioff M, Moulin H, Daskalakis C, eds. *Proc. 18th ACM Conf. on Econom. and Comput.* (ACM, New York), 629–646.
- Conitzer V, Freeman R, Shah N, Vaughan J (2019) Group fairness for the allocation of indivisible goods. Hentenryck PV, Zhou Z-H, eds. *Proc. 33rd AAAI Conf. on Artificial Intelligence* (AAAI Press, Palo Alto, CA), 1853–1860.
- Dubins L, Spanier E (1961) How to cut a cake fairly. *Amer. Math. Monthly* 68(1):1–17.
- Foley D (1967) Resource allocation and the public sector. *Yale Econom. Essays* 7:45–98.
- Freeman R, Shah N, Vaish R (2020) Best of both worlds: Ex-ante and ex-post fairness in resource allocation. Ostrovsky M, Procaccia A, Biro P, Hartline J, eds. *Proc. 21st ACM Conf. on Econom. and Comput.* (ACM, New York), 21–22.
- Freeman R, Sikdar S, Vaish R, Xia L (2019) Equitable allocations of indivisible goods. Kraus S, ed. *Proc. 28th Internat. Joint Conf. on Artificial Intelligence* (International Joint Conferences on Artificial Intelligence, California), 280–286.
- Gajdos T, Tallon J (2002) Fairness under uncertainty. *Econom. Bull.* 4(18):1–7.
- Gamow G, Stern M (1958) *Puzzle-Math* (Viking Press).
- Gourvès L, Monnot J, Tlilane L (2014) Near fairness in matroids. Schaub T, Friedrich G, O’Sullivan B, eds. *Proc. 21st Eur. Conf. on Artificial Intelligence* (IOS Press, Amsterdam), 393–398.
- Halpern D, Procaccia AD, Psomas A, Shah N (2020) Fair division with binary valuations: One rule to rule them all. Chen X, Gravin N, Hoefer M, Mehta R, eds. *Proc. 16th Internat. Conf. on Web and Internet Econom.* (Springer, Berlin), 370–383.
- Hylland A, Zeckhauser R (1979) The efficient allocation of individuals to positions. *J. Political Econom.* 87(2):293–314.
- Johnson DM, Dulmage A, Mendelsohn N (1960) On an algorithm of G. Birkhoff concerning doubly stochastic matrices. *Canadian Math. Bull.* 3(3):237–242.
- Katta A-K, Sethuraman J (2006) A solution to the random assignment problem on the full preference domain. *J. Econom. Theory* 131(1):231–250.
- Kojima F (2009) Random assignment of multiple indivisible objects. *Math. Social Sci.* 57(1):134–142.
- Lipton RJ, Markakis E, Mossel E, Saberi A (2004) On approximately fair allocations of indivisible goods. Breese J, Feigenbaum J, Seltzer M, eds. *Proc. 6th ACM Conf. on Electronic Commerce* (ACM, New York), 125–131.
- Matoušek J, Björner A, Ziegler GM, et al. (2003) *Using the Borsuk-Ulam Theorem: Lectures on Topological Methods in Combinatorics and Geometry* (Springer, Berlin).
- Moulin H (2003) *Fair Division and Collective Welfare* (MIT Press, Cambridge, MA).
- Nesterov A (2017) Fairness and efficiency in strategy-proof object allocation mechanisms. *J. Econom. Theory* 170:145–168.
- Orlin J (2010) Improved algorithms for computing Fisher’s market clearing prices. Schulman LJ, ed. *Proc. 42nd ACM Sympos. on Theory of Comput.* (ACM, New York), 291–300.
- Plummer MD, Lovász L (1986) *Matching Theory* (Elsevier, New York).
- Steinhaus H (1948) The problem of fair division. *Econometrica* 16:101–104.
- Varian H (1974) Equity, envy and efficiency. *J. Econom. Theory* 9:63–91.
- Végh L (2016) A strongly polynomial algorithm for a class of minimum-cost flow problems with separable convex objectives. *SIAM J. Comput.* 45(5):1729–1761.
- von Neumann J (1953) A certain zero-sum two-person game equivalent to the optimal assignment problem. Kuhn W, Tucker AW, eds. *Contributions to the Theory of Games*, vol. 2 (Princeton University Press, Princeton, NJ), 5–12.

Haris Aziz is a Scientia associate professor at University of New South Wales and the director of the University of New South Wales Artificial Intelligence Institute. His research interests lie at the intersection of artificial intelligence, theoretical computer science, and mathematical social sciences.

Rupert Freeman is an assistant professor of business administration in the quantitative analysis area at the University of Virginia’s Darden School of Business. He is interested in game theory and mechanism design with applications to resource allocation, forecasting, and social choice.

Nisarg Shah is an assistant professor of computer science at the University of Toronto, a research lead for Ethics of Artificial Intelligence at the Schwartz Reisman Institute for Technology and Society, and a faculty affiliate of the Vector Institute for Artificial Intelligence. He works on theoretical foundations of artificial intelligence, multiagent systems, and algorithmic fairness.

Rohit Vaish is an assistant professor in the Department of Computer Science and Engineering at the Indian Institute of Technology Delhi. His research is in problems at the interface of computer science and economics, specifically in the area of computational social choice, spanning topics such as voting, matching, and fair division.