
Fair and Efficient Allocations Without Obvious Manipulations

Alexandros Psomas
Purdue University
apsomas@cs.purdue.edu

Paritosh Verma
Purdue University
verma136@purdue.edu

Abstract

We consider the fundamental problem of allocating a set of indivisible goods among strategic agents with additive valuation functions. It is well known that, in the absence of monetary transfers, Pareto efficient and truthful rules are dictatorial, while there is no deterministic truthful mechanism that allocates all items and achieves envy-freeness up to one item (EF1), even for the case of two agents. In this paper, we investigate the interplay of fairness and efficiency under a relaxation of truthfulness called non-obvious manipulability (NOM), recently proposed by [TM20]. We show that this relaxation allows us to bypass the aforementioned negative results in a very strong sense. Specifically, we prove that there are deterministic and EF1 algorithms that are not obviously manipulable, and the algorithm that maximizes utilitarian social welfare (the sum of agents’ utilities), which is Pareto efficient but not dictatorial, is not obviously manipulable for $n \geq 3$ agents (but obviously manipulable for $n = 2$ agents). At the same time, maximizing the egalitarian social welfare (the minimum of agents’ utilities) or the Nash social welfare (the product of agents’ utilities) is obviously manipulable for any number of agents and items. Our main result is an approximation preserving black-box reduction from the problem of designing EF1 and NOM mechanisms to the problem of designing EF1 algorithms. En route, we prove an interesting structural result about EF1 allocations, as well as new “best-of-both-worlds” results (for the problem without incentives), that might be of independent interest.

1 Introduction

We consider the fundamental problem of allocating a set of indivisible items among strategic agents with additive preferences. It is well-understood that, in the absence of monetary transfers, fairness, efficiency and truthfulness cannot be reconciled, in a very strong sense. For example, a serial dictatorship (arguably the most unfair rule) is the *unique* truthful and Pareto efficient mechanism [KM02, EK03], even for the case of two agents and randomized mechanisms (or, equivalently, divisible items) [Sch96]. At the same time, achieving envy-freeness up to one item (EF1) [Bud11], a popular fairness notion for indivisible goods that is compatible with Pareto efficiency [CKM⁺19], is impossible for truthful and deterministic mechanisms to achieve, even for two additive agents [ABCM17].

The standard abstraction for a strategic agent assumes that, when presented with an algorithm, the agent will perfectly understand it and optimally respond to it. Therefore, a truthful mechanism should protect against all deviating strategies. This modeling is overly pessimistic since, for example, in practice agents do not have exact knowledge about others’ preferences. Since truthfulness rules out reasonable solutions, previous work typically ignores strategic issues altogether. In this paper we aim to escape the aforementioned impossibility results by relaxing truthfulness: our goal is to design mechanisms that protect only against *obviously dominant* deviating strategies.

Informally, a strategy is obviously dominant if it guarantees an outcome better than any outcome of every other strategy [Li17]. Designing obviously truthful mechanisms, those where honest reporting is obviously dominant, is a problem that has attracted significant attention in recent years [AG18, KV19, FV19, AMN20]. Since dishonesty is also a strategy, a deviation can be profitable but not obviously profitable. A manipulation is obvious if it yields a higher utility than truth-telling in either the best or worst-case scenarios. Seminal work of [TM20] defines this notion, and shows that relaxing truthfulness to non-obvious manipulability can strictly improve the designer’s objective in a number of domains (e.g. two-sided matching). In this paper we aim to provide a comprehensive study of what can and what cannot be achieved by non-obviously manipulable mechanisms in the fundamental domain of fair and efficient allocation of indivisible goods.

Our Contribution. We first consider the problem of designing a non-obviously manipulable mechanism (NOM) which always outputs allocations that are envy-free up to one item (EF1), a task that is impossible for truthful mechanisms to achieve [ABCM17]. We prove that the Round-Robin procedure (agents choose items one at a time, following a fixed order) which is known to always output EF1 allocations, is also not obviously manipulable (Theorem 1), giving a separation between truthfulness and non-obviously manipulability for our problem. In fact, we show that NOM is compatible with even stronger notions of fairness. Specifically, a randomized algorithm can be ex-post EF1 (i.e. the output allocation is always EF1), and simultaneously be envy-free in expectation (ex-ante EF). There are known algorithms that satisfy this “best-of-both-worlds” guarantee [FSV20, Azi20b]; we show that the PS-Lottery algorithm of [Azi20b] is not-obviously manipulable. This result uses a connection between ex-ante proportionality and not obvious manipulability; we later exploit this connection to establish new negative “best-of-both-worlds” results.

We proceed to study efficient algorithms. We focus on the three most prominent notions of efficiency: utilitarian social welfare (the sum of agents’ utilities), egalitarian social welfare (the minimum of agents’ utilities), and Nash social welfare (the geometric mean of agents’ utilities). For the case of $n = 2$ agents all three notions are incompatible with non obvious manipulability, i.e. every (deterministic or randomized) algorithm that always outputs an integral solution that is optimal with respect to any of these objectives is obviously manipulable (Theorem 2). This is also true for the case of more than three agents for egalitarian and Nash social welfare (Theorems 8 and 4). The high level intuition for these impossibilities is as follows. There are instances where a specific agent¹ gets only her least favorite item, in the worst-case. However, maximizing egalitarian or Nash social welfare ensures that the number of agents that get non-zero utility are maximized. Therefore, by reporting that she only values a single item (her favorite item) to the mechanism, this agent can force the mechanism to give her this item (or increase the probability that she gets it) in the worst-case. Note that lack of agents’ knowledge about each others’ preferences is used as a justification for why manipulations might not be a first-order concern in our domain [CKM⁺19]; the fact that maximizing Nash social welfare is manipulable even when agents know nothing about each other further highlights the need for more nuanced formal models between “agents are always honest” and “agents are perfectly rational and all-knowing expected utility maximizers.” Surprisingly, the same is not true for utilitarian social welfare: there is a NOM mechanism that always outputs an allocation that maximizes utilitarian social welfare for the case of $n \geq 3$ agents (Theorem 3). This gives an efficient and not obviously manipulable mechanism that is not dictatorial.

Since maximizing Nash social welfare, an objective that simultaneously achieves Pareto efficiency and EF1, is obviously manipulable, the next natural question is whether there are any other Pareto efficient and EF1 algorithms that satisfy NOM. Our main result answers this question in the affirmative. In fact, we prove that the problem of designing a Pareto efficient, EF1 and NOM *mechanism* is exactly as hard as designing a Pareto efficient and EF1 *algorithm*. We give a black-box reduction, that preserves Pareto efficiency guarantees, which given an algorithm that always outputs (clean and non-wasteful) EF1 allocations, outputs a mechanism that is not obviously manipulable and always outputs EF1 allocations. There are two crucial steps in our reduction. Given the valuation function \mathbf{v}_i of agent i , let $\text{EF1}(i, \mathbf{v}_i)$ be the set of allocations that are clean, non-wasteful and EF1 for agent i . First, our reduction ensures that, when agent i reports her valuation to be \mathbf{b}_i , every single allocation in $\text{EF1}(i, \mathbf{b}_i)$ is possible, i.e. for every allocation $A \in \text{EF1}(i, \mathbf{b}_i)$ there are reports for the other agents such that A is the output of the mechanism. Second, we prove the following intuitive structural result, that might be of independent interest. For an agent i having valuation function \mathbf{v}_i , the worst-case allocation in $\text{EF1}(i, \mathbf{v}_i)$ is better than the worst-case allocation in $\text{EF1}(i, \mathbf{b}_i)$ (Lemma 4). These facts

¹Informally, the agent that the tie-breaking rule favors.

Table 1: Existence of NOM mechanisms for various fairness and efficiency notions.

Fairness & Efficiency Notions	Existence of NOM Mechanism
EF1	✓ [Theorem 1]
Ex-ante EF & Ex-post EF1	✓ [Theorem 7]
Utilitarian Social Welfare	✗ $n = 2$ agents [Theorem 2] ✓ $n \geq 3$ agents [Theorem 3]
Egalitarian Social Welfare	✗ [Theorem 8]
Nash Social Welfare	✗ [Theorem 4]
fPO + EF1	✓ [Theorem 5]
Ex-ante EF & Ex-post fPO + EF1	✗ [FSV20]

combined establish the “worst-case” part from the definition of NOM, which is the most challenging step. As a direct application of our reduction, by giving as input the algorithmic results [GM21], we get a fractionally Pareto efficient, EF1 and NOM mechanism for additive agents; this last result cannot be improved by adding ex-ante fairness guarantees (due to a theorem of [FSV20]). Finally, despite the aforementioned technical subtleties, the “code” of our reduction is fairly simple: it checks whether some allocations are EF1, and if not it calls the black-box algorithm. This simplicity is a feature, with concrete practical implications: by adding a few lines of code to the implementation of any EF1 and Pareto efficient algorithm (e.g. the algorithm that maximizes Nash social welfare, which is used by the popular website Spliddit [GP15]) one can maintain these guarantees, while provably protecting against certain types of manipulations.

Our results show a connection between certain notions of fairness and not obvious manipulability: positive algorithmic results for EF1 and “best-of-both-worlds” guarantees can be used to get positive results for not obviously manipulable mechanisms. In Appendix E we exploit this connection in the other direction, and show how negative results for not obviously manipulable mechanisms can be used to prove new negative results for “best-of-both-worlds” algorithms. Specifically, we prove that it is impossible to achieve ex-ante proportionality, ex-post Pareto efficiency while ex-post maximizing the number of agents with positive utility. As direct corollaries we recover a known result of [FSV20] that ex-ante proportional and ex-post MNW allocations do not exist, as well as new results: it is impossible to achieve ex-ante proportionality by randomizing over allocations that are (1) ex-post Pareto efficient and ex-post egalitarian, or (2) ex-post leximin.

Related work. [TM20] define the notion of not obvious manipulability. They show that a number of known mechanisms that are not truthful satisfy not obviously manipulability, e.g. stable mechanisms in the context of two-sided matching are not obviously manipulable (in direct contrast to [Rot82], who shows that there exists no mechanism that is both stable and truthful). At the same time, many known mechanisms that are not truthful are obviously manipulable, e.g. the first-price auction, and more generally the pay-as-bid auction.

The notion of non obvious manipulability has been explored in some very recent works [OSH19, AL21]. [OSH19] study cake-cutting and show that, as opposed to truthfulness, NOM is compatible with proportionality: an adaptation of the well-known moving knife procedure satisfies both properties. They also observe that every proportional (direct revelation) rule satisfies the worst-case guarantee required for NOM (we show a similar statement in Lemma 5), while every Pareto efficient (direct revelation) rule satisfies the best-case guarantee required for NOM. [AL21] explore obvious manipulations in voting theory. They give sufficient conditions for voting rules to be not obviously manipulable, and show that a number of voting rules escape the pessimistic conclusions of the Gibbard-Satterthwaite theorem, e.g., the Borda rule is not obviously manipulable.

Other than relaxing the notion of incentive compatibility, one can escape the aforementioned impossibility results by restricting agents’ valuations, e.g. focus on dichotomous [HPPS20, BEF21, BCIZ21, BV21] or Leontief valuations [GZH⁺11, FGP14, PPS15], or by using money-burning (wasting resources) as a substitute for payments [HR08, CGG13, FTTZ16, FGPS19, ACGH20].

2 Preliminaries

We consider the problem of allocating a set M of m items among a set N of n agents with additive utilities. We use $[k]$ to denote the set $\{1, 2, \dots, k\}$ for any positive integer $k \in \mathbb{Z}_+$. Proofs pertaining to randomized mechanisms are deferred to the appendix; see Appendix A for the corresponding preliminaries on randomized allocations/mechanisms.

Allocations. A *fractional allocation* $A \in [0, 1]^{n \cdot m}$ is a $m \times n$ matrix that defines for each agent $i \in N$ and item $j \in M$ the fraction $A_{i,j}$ of the item j that the agent i receives. We represent a fractional allocation as $A = (A_1, \dots, A_n)$ where $A_i = (A_{i,1}, \dots, A_{i,m}) \in [0, 1]^m$ denotes the fractions of all items allocated to agent i . A feasible allocation satisfies $\sum_{i \in N} A_{i,j} \leq 1$, for all $j \in M$. A fractional allocation A is *integral* if $A_i \in \{0, 1\}^m$ for all agents $i \in N$. An integral allocation can be equivalently defined as n disjoint subsets of the set of items M , this representation is often convenient. Let $\Pi_n(M) = \{(S_1, \dots, S_n) \mid \cup_{i=1}^n S_i \subseteq M \text{ and } \forall i \neq j, S_i \cap S_j = \emptyset\}$ denote the set of all n ordered disjoint subsets of the set of items M . Given an integral allocation $A = (A_1, \dots, A_n)$, we can interpret the binary vectors $A_i = (A_{i,1}, \dots, A_{i,m})$ as subsets of items $A_i := \{j \in M \mid A_{i,j} = 1\}$. This allows us to view an integral allocation A as n ordered disjoint subsets of items, i.e., $A \in \Pi_n(M)$. An integral allocation A is *complete* if $\cup_{i=1}^n A_i = M$ and *partial* if $\cup_{i=1}^n A_i \subset M$. Unless stated otherwise, we will use the term allocation to refer to complete allocations. We will use the term *bundle* to refer to any subset of items.

Preferences. Each agent $i \in N$ has a private valuation function $v_i(\cdot)$ that outputs the utility that agent i derives from a given set (or fractions) of items. We use the notation $v_i(X_i)$ (resp. $v_i(A_i)$) to denote the utility that agent i gets from the items allocated to her in a fractional allocation X (resp. integral allocation A). We consider agents with additive utilities. An additive agent $i \in N$ has a non-negative valuation $v_{i,j}$ for receiving the entirety of item j ; her utility for an allocation X is $v_i(X_i) = \sum_{j \in M} X_{i,j} v_{i,j}$; for an integral allocation A , the utility is simply $v_i(A_i) = \sum_{j \in A_i} v_{i,j}$.

Mechanisms. A mechanism \mathcal{M} elicits “bids” (i.e. reported valuations) $\mathbf{b} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ from every agent $i \in N$, and outputs a feasible allocation. We write $X_{i,j}(\mathbf{b})$ for the fraction of item j allocated to agent i when each agent $j \in N$ reports a valuation \mathbf{b}_j . We use the notation \mathbf{b} (and \mathbf{b}_i) to refer to the input to a mechanism, and \mathbf{v} (and \mathbf{v}_i) to refer to the true valuations of agents. A mechanism \mathcal{M} is *deterministic* if for every reported valuations \mathbf{b} it deterministically outputs an integral allocation. We will use $\mathcal{M}(\mathbf{b}) = (\mathcal{M}_1(\mathbf{b}), \mathcal{M}_2(\mathbf{b}), \dots, \mathcal{M}_n(\mathbf{b}))$ to denote the integral allocation that a mechanism \mathcal{M} outputs given reported valuations $\mathbf{b} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ as input, here $\mathcal{M}_i(\mathbf{b}) \subseteq M$ represents the bundle of items that agent i receives in the allocation $\mathcal{M}(\mathbf{b})$. For a deterministic mechanism \mathcal{M} , the value $v_i(\mathcal{M}_i(\mathbf{b}))$ denotes the utility of agent i for the allocation output by \mathcal{M} on input \mathbf{b} .

Notions of incentive compatibility. A mechanism \mathcal{M} is *truthful* if agents cannot strictly improve their utility by misreporting their valuation, i.e. for all $i \in N$, valuations $\mathbf{v}_i, \mathbf{b}_i$, and reports of the other agents \mathbf{v}_{-i} , $v_i(\mathcal{M}_i(\mathbf{v}_i, \mathbf{v}_{-i})) \geq v_i(\mathcal{M}_i(\mathbf{b}_i, \mathbf{v}_{-i}))$. Our work focuses on a notion of incentive compatibility that is a relaxation of truthfulness called *not obvious manipulability*.

Definition 1 (Not Obviously Manipulable [TM20]). A mechanism \mathcal{M} is *not obviously manipulable* (NOM) if for every agent $i \in N$ with valuation function \mathbf{v}_i , and every possible report \mathbf{b}_i of agent i , the following two inequalities hold:

- (1) $\min_{\mathbf{v}_{-i}} v_i(\mathcal{M}_i(\mathbf{v}_i, \mathbf{v}_{-i})) \geq \min_{\mathbf{v}_{-i}} v_i(\mathcal{M}_i(\mathbf{b}_i, \mathbf{v}_{-i}))$.
- (2) $\max_{\mathbf{v}_{-i}} v_i(\mathcal{M}_i(\mathbf{v}_i, \mathbf{v}_{-i})) \geq \max_{\mathbf{v}_{-i}} v_i(\mathcal{M}_i(\mathbf{b}_i, \mathbf{v}_{-i}))$.

Intuitively, if a mechanism is NOM then an agent cannot increase her worst-case utility or her best-case utility (computed with respect to the true valuation) by misreporting her valuation. If either the worst-case or the best-case utility can be improved then the mechanism is *obviously manipulable*.

Notions of efficiency. An integral allocation $A = (A_1, \dots, A_n)$ is *Pareto efficient* (or PO) iff there is no integral allocation $A' = (A'_1, \dots, A'_n)$ such that for all agents $i \in N$, $v_i(A'_i) \geq v_i(A_i)$, and for at least one agent this inequality is strict. An (integral or fractional) allocation $X = (X_1, \dots, X_n)$ is *fractionally Pareto efficient* (or fPO) iff there is no fractional allocation $X' = (X'_1, \dots, X'_n)$ such that for all agents $i \in N$, $v_i(X'_i) \geq v_i(X_i)$, and for at least one agent this inequality is strict. Note that fractional Pareto efficiency is a strictly stronger notion than Pareto efficiency. An (integral or fractional) allocation X is α -approximately (resp. fractionally) Pareto efficient if there is no integral

allocation (resp. fractional allocation) $X' = (X'_1, X'_2, \dots, X'_n)$ such that $\alpha \cdot \mathbf{v}_i(X'_i) \geq \mathbf{v}_i(X_i)$, with at least one of these inequalities strict [RF90, ILWM17, FGPS19, ZP20]. Notice that for $\alpha = 1$ we exactly recover Pareto efficiency (resp. fractional Pareto efficiency).

A (partial) allocation $A = (A_1, A_2, \dots, A_n)$ is *non-wasteful* iff for each $i \in \mathbb{N}$, $v_{i,j} = 0$ for every unallocated item $j \in \mathbb{M} \setminus \cup_{k \in \mathbb{N}} A_k$ ². A bundle $S \subseteq \mathbb{M}$ is *clean* wrt valuation \mathbf{v}_i if the items comprising S have positive value, i.e., $\mathbf{v}_i(g) > 0$ for all $g \in S$. Further, a (partial) allocation $A = (A_1, A_2, \dots, A_n)$ is clean if for each $i \in \mathbb{N}$ the bundle A_i is clean wrt valuation \mathbf{v}_i .

Often, we are interested in computing or approximating specific points (i.e., allocations) present on the Pareto frontier that have additional desirable properties. The *utilitarian social welfare* of an allocation X , denoted by $\text{SW}(X)$, is defined as the sum of utilities that each agent gets in the allocation X , i.e., $\text{SW}(X) = \sum_{i \in \mathbb{N}} \mathbf{v}_i(X_i)$. The *Nash social welfare* of an (integral or fractional) allocation X , denoted by $\text{NSW}(X)$, is defined as the geometric mean of the utilities of agents in the allocation X , i.e., $\text{NSW}(X) = (\prod_{i \in \mathbb{N}} \mathbf{v}_i(x_i))^{\frac{1}{n}}$. An *integral* allocation that maximizes the Nash social welfare, among all integral allocations, is called a *Nash social welfare maximizing* (or MNW) allocation. There are instances where every integral allocation A is such that $\text{NSW}(A) = 0$, i.e., there is always an agent having zero utility. To cover such an edge case, integral MNW allocations are defined as follows [CKM⁺19]. An integral allocation is MNW iff (i) it maximizes, among the set of all integral allocations, the number of agents having positive utility and (ii) for any such maximal set of agents S , it maximizes the geometric mean of the utilities of agents in S .

Notions of fairness. An (integral or fractional) allocation X is called *envy-free* (EF) if for every pair of agents $i, j \in \mathbb{N}$, agent i values her allocation at least as much as the allocation of agent j , i.e., $\mathbf{v}_i(X_i) \geq \mathbf{v}_i(X_j)$. Achieving envy-freeness is impossible for integral allocations (consider the case of a single item and two agents that both want it), so relaxation of envy-freeness are considered. An integral allocation A is *envy-free up to one item* (EF1) if for every pair of agents $i, j \in \mathbb{N}$, where $A_j \neq \emptyset$, agent i values her allocation at least as much as the allocation of agent j , subject to the removal of one item from agent j 's bundle, i.e., $\mathbf{v}_i(A_i) \geq \mathbf{v}_i(A_j \setminus \{g\})$ for some item $g \in A_j$.

Finally, we are often interested in, so called, best-of-both-worlds guarantees. Let \mathcal{P} be a fairness or efficiency notion for integral allocations and \mathcal{Q} be a fairness or efficiency notion for fractional allocations. A randomized allocation \mathbf{R} (see Appendix A for definition), which outputs an integral allocation A^z with probability p^z , satisfies the notion \mathcal{P} *ex-post* if each integral allocation A^z in the support of \mathbf{R} satisfies \mathcal{P} . Additionally, the randomized allocation \mathbf{R} satisfies the notion \mathcal{Q} *ex-ante* if the expected fractional allocation $Y = \sum_{z=1}^k p^z \cdot A^z$ corresponding to \mathbf{R} satisfies the notion \mathcal{Q} .

3 Fair Mechanisms

In this section we study whether non obvious manipulability is compatible with EF1. For deterministic mechanisms, it is known that no truthful mechanism can always outputs an EF1 allocation, even for the case of two agents [ABCM17]. In sharp contrast, we show that Round-Robin, arguably the simplest EF1 algorithm, is not obviously manipulable. Recall that the Round-Robin algorithm allocates items to agents over a sequence of rounds. In each round, the agents choose one item each (a highest-value remaining item, as per their valuation³) following a fixed, arbitrary order.

Theorem 1. *Round-Robin is not obviously manipulable.*

Proof. We prove that the two inequalities in Definition 1 hold for Round-Robin. Let $i \in \mathbb{N}$ be the i -th agent in the Round-Robin order, and let \mathbf{v}_i be her valuation vector. Assume without loss of generality that $v_{i,j} \geq v_{i,j+1}$ for all $j = 0, \dots, m-1$. Let ℓ be the number of items agent i receives (either $\lfloor \frac{n}{m} \rfloor$ or $\lfloor \frac{n}{m} \rfloor + 1$, depending on i, n and m), and notice that ℓ does not depend on her report.

We first prove inequality (1) from the definition of NOM, i.e. the worst-case guarantee. The k -th time agent i gets to pick an item, there is an un-allocated item that she values at least $v_{i,i+(k-1)n}$, since only $(k-1)n + i - 1$ items have been allocated at that point. Therefore, the worst-case for agent i under honest reporting is realized when all other agents rank the items in the same order as her, and her utility is exactly $\sum_{k=1}^{\ell} v_{i,i+(k-1)n}$. Now consider the case when all agents other than i

²Note that, every Pareto efficient (fractionally Pareto efficient) integral allocation is non-wasteful.

³Ties can be broken in an arbitrary manner.

rank items in the same order as i (i.e. $v_{i',j} \geq v_{i',j+1}$ for all $i' \in N$), and i reports some vector \mathbf{b}_i . Let j_1, \dots, j_ℓ be the items she receives, in the order she picked them. First, without loss of generality, i picks these items in decreasing order with respect to her true valuation, i.e. $v_{i,j_k} \geq v_{i,j_{k+1}}$: if she picks an item k at round t , and at a later round t' she could pick an item k' with $v_{i,k'} > v_{i,k}$, then by the choice of valuation for agents other than i , picking k' in round t and k in round t' is also possible. Second, $v_{i,j_k} \leq v_{i,i+(k-1)n}$: when i picks an item for the k -th time, $(i-1) + (k-1)(n-1)$ of her top $i + (k-1)n$ items have been picked by other agents, and $k-1$ of her top $i + (k-1)n$ items have been picked by herself (since items are picked in decreasing order with respect to the true valuation). Therefore $v_{i,i+(k-1)n}$ is the largest value v_{i,j_k} can have; i 's utility is at most $\sum_{k=1}^{\ell} v_{i,i+(k-1)n}$. Thus, the worst-case utility is maximized under truthful reporting.

Next, we prove inequality (2) (the best-case guarantee). Since agent i gets exactly ℓ items, her utility is at most $\sum_{k=1}^{\ell} v_{i,k}$. This utility is realized when she reports \mathbf{v}_i , and everyone else ranks the items in the opposite order, i.e. $v_{i',j} < v_{i',j+1}$ for all i' , and it cannot be improved upon. \square

In fact, we can strengthen this result to give a “best-of-both-worlds” NOM mechanism. Specifically, we prove that the PS-Lottery algorithm of [Azi20b], which outputs randomized allocations that are ex-ante EF and ex-post EF1, is NOM in expectation; proof has been deferred to Appendix B.

4 Efficient Mechanisms

In this section we study whether natural efficiency notions are compatible with NOM. We consider the three most popular notions of efficiency: utilitarian social welfare, Nash social welfare and egalitarian social welfare (see Appendix C). Under truthfulness, the *only* Pareto efficient and truthful algorithm is a dictatorship, which immediately implies that one cannot truthfully achieve any non-trivial approximations with respect to any of these notions. Omitted proofs can be found in Appendix C.

Utilitarian Social Welfare. We start by considering the utilitarian social welfare maximizing algorithm, i.e. the algorithm which allocates each item to the agent that values it the most. When the winner for an item is not unique the algorithm needs to break ties; the choice of the tie-breaking rule will be crucial for our positive result in this section. In the context of fair division it is standard to assume that agents’ values are normalized when analyzing utilitarian social welfare. Specifically, the most common assumption is that the agents’ values add up to 1; see [Azi20a] for a number of justifications for this assumption. For the remainder of this section we also assume that $\sum_{j \in M} v_{i,j} = 1$ for all agents i .⁴ We note that without this assumption, the utilitarian social welfare maximizing algorithm is NOM, since no matter what an agent reports, the best case for her is that she wins all items she values positively, and the worst case is that she loses all items, so no misreport can help increase either the worst-case or best-case utility.

Theorem 2. *Every (randomized or deterministic) mechanism for $n = 2$ agents that always outputs utilitarian social welfare maximizing allocations is obviously manipulable.*

Surprisingly, this impossibility result does not hold for more than two agents.

Theorem 3. *The utilitarian social welfare maximizing algorithm, coupled with an appropriate tie-breaking rule, is not obviously manipulable for $n \geq 3$ agents.*

Proof. First, we describe our algorithm. The algorithm, given reports $\mathbf{b}_1, \dots, \mathbf{b}_n$, allocates each item to an agent with the largest reported value. In case of a tie, the item is allocated to the agent with the smallest index (i.e. $1 > 2 > 3 \dots$), *except* if the tie is exactly between agents 1 and n , in which case the item goes to agent n . Consider an arbitrary agent i with true valuation \mathbf{v}_i .

Towards proving inequality (1), we have that if $b_{i,j} < 1$ for all $j \in M$, there exists a choice for \mathbf{v}_{-i} such that agent i does not have the highest value for any item (e.g., some agent k could have a value of 1 for i 's favorite item, and a different agent k' can out-bid i in all remaining items). If $b_{i,j} = 1$ for some $j \in M$, then, again, i can again end up with no items, since by the choice of our tie-breaking rule, all agents can lose in the tie-breaking (agents other than 1 lose to some smaller index agent,

⁴Equivalently, we can define the utilitarian social welfare maximizing algorithm to be the aforementioned algorithm executed on transformed valuations where $\hat{v}_{i,j} = \frac{v_{i,j}}{\sum_{j \in M} v_{i,j}}$ for all $i \in N$.

while agent 1 loses to agent n). Therefore, no matter what agent i reports, in the worst-case she gets a utility of zero. Next we prove inequality (2), the best-case guarantee. By the normalization assumption and tie-breaking rule, any agent i (except agent 1) cannot out-bid every other agent on every single item, no matter what \mathbf{b}_i and \mathbf{v}_{-i} are. Therefore, the best-case outcome for agent i is that she is allocated all items except her least favorite one. Additionally, note that this outcome can be realized when agent i reports honestly and \mathbf{v}_{-i} is such that $v_{i',j} = 1$, for all $i' \neq i$, where j is the item that i values the least (i.e. $j \in \operatorname{argmin}_{k \in M} \{v_{i,k}\}$). Finally, if agent 1 reports her true valuation, the best-case happens when all other agents also report agent 1's true valuation — in which case, as per the tie-breaking rule, she gets all the items; this outcome also cannot be improved upon. \square

Nash Social Welfare. The next natural maximization objective, and arguably the most popular one in the context of fair division, is the Nash social welfare. Here, we show that — unlike utilitarian social welfare maximization — there does not exist a NOM mechanism that always outputs allocations that maximize the Nash social welfare, for any number of agents. The case of $n = 2$ agents follows from Theorem 10 (Appendix E) that details our “best-of-both-worlds” results. In the following theorem (proof deferred to Appendix C), we show that the impossibility extends to $n \geq 3$ agents.

Theorem 4. *Every (randomized or deterministic) mechanism that always outputs an allocation that maximizes the Nash social welfare is obviously manipulable, even for 3 additive agents and 4 items.*

5 Fair and Efficient Mechanisms

Maximizing the Nash social welfare results in allocations that are fair in addition to being economically efficient [CKM⁺19]. However, as we established in Theorem 4, mechanisms that output MNW allocations are obviously manipulable for any number of agents.

In this section we state our main result, where we show that there is a deterministic NOM mechanism that is fair and economically efficient for additive valuations. This result is established by showing a black-box reduction from the problem of designing NOM mechanisms that output EF1 allocations to the problem of designing algorithms that output EF1 allocations. Additionally, this black-box reduction preserves Pareto efficiency guarantees. Missing proofs can be found in Appendix D.

The following theorem formally states our main result.

Theorem 5. *For additive valuations, there exists a black-box reduction, which preserves Pareto efficiency guarantees, from the problem of designing a NOM and EF1 mechanism, to designing an algorithm that computes clean, non-wasteful and EF1 allocations.*

By combining Theorem 5 with known algorithmic results we can get NOM mechanisms, with the same fairness and efficiency guarantees. Specifically, [GM21] give a pseudo-polynomial time algorithm that computes fractionally Pareto efficient and EF1 allocations. Fractional Pareto efficiency implies non-wastefulness, and without loss of generality we can assume that this algorithm outputs clean allocations.⁵ Hence, by applying Theorem 5 we get the following application.

Corollary 1. (via [GM21]). *For agents with additive valuations, there exists a fractionally Pareto efficient, EF1, and NOM mechanism, that runs in pseudo-polynomial time.*

The Reduction. Our reduction, Mechanism 1, takes as input reported valuations $\mathbf{b} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$, and black-box access to a deterministic algorithm \mathcal{M}^* .⁶ Our reduction requires the algorithm \mathcal{M}^* to always output non-wasteful, clean EF1 allocations for every possible input valuation functions.

The reduction is based on two key ideas, first, through a careful construction of cases we ensure that if an agent $i \in N$ reports valuation \mathbf{b}_i , then the set of allocations that can be returned by our reduction, Mechanism 1, for every possible \mathbf{b}_{-i} can be precisely characterized (Lemma 3). Second, we prove a structural result (Lemma 4) concerning this set of possible output allocation. This structural result plays a central role in establishing that Mechanism 1 is not obviously manipulable. Further, we show that, by construction, our reduction always outputs fPO + EF1 allocations.

We begin by defining some notation. For each agent $i \in N$, let D_i be the set of goods that have strictly positive value for i i.e. $D_i := \{g \in M \mid b_{i,g} > 0\}$. Let $\widehat{D}_i := M \setminus \cup_{j \neq i} D_j$ be the goods

⁵Starting from an fPO and EF1 allocation A , making each bundle A_i clean preserves both properties.

⁶Algorithm \mathcal{M}^* could possibly be computationally inefficient.

remaining after removing all goods desired by agents other than i . Let R_i , for each agent $i \in N$, be the indicator for the event that the subsets $\{D_j\}_{j \in N} \setminus \{D_i\}$ are pairwise disjoint, i.e. $R_i = 1$ iff $(\{D_j\}_{j \in N} \setminus \{D_i\})$ are pairwise disjoint, and $R_i = 0$ otherwise.

Mechanism 1 sequentially considers four cases based on the sets $\{D_j\}_{j=1}^n$ and the values $\{R_j\}_{j=1}^n$, to find a temporary allocation A^* . See Appendix D.2 for the pseudo-code.

Case I: The sets $\{D_j\}_{j=1}^n$ are pairwise disjoint (equivalently, $R_i = 1$ for all $i \in N$). In this case, A^* allocates the bundle D_j to agent j for each agent $j \in N$.

Case II: $R_i = 1$ for exactly one agent $i \in N$. This can occur if D_i intersects two or more D_j s (and these are the only intersections among pairs of subsets in $\{D_k\}_{k=1}^n$). If allocating the bundle \widehat{D}_i to agent i , and the bundle D_j to each agent $j \in N$, for each $j \neq i$ results in an EF1 allocation, then A^* is set to this allocation. Otherwise, if this allocation is not EF1, A^* is the allocation returned by \mathcal{M}^* .

Case III: There are exactly two agents $i, j \in N$ such that $R_i = R_j = 1$. The only way this is possible is if D_i, D_j intersect each other and any other pair of subsets D_k, D_l where $\{k, l\} \neq \{i, j\}$, are disjoint. In this case, Mechanism 1 considers whether the set of goods $D_i \cap D_j$ are valued more by agent i or agent j ; each of these two subcases are similar to *Case II* (see Lines 12-21).

Case IV: None of the previous cases holds (equivalently, $R_i = 0$ for all $i \in N$). In this case, A^* is simply the allocation returned by \mathcal{M}^* .

The last step of Mechanism 1 (Line 24) is to make the bundle allocated to each agent i in the temporary allocation *clean*, i.e. remove items from her bundle that she does not value. This is necessary for one of our technical lemmas (specifically, for characterizing the set of all allocations that are possible outputs of our mechanism). Note that this step does not affect efficiency or envy-freeness up to one item (i.e. if the allocation satisfied any of these notions before this step, it continues to do so).

For the case of two agents, our reduction is especially simple: give each agent i all items she has a positive value and the other agent has a zero value, and give all items that both agents want (if any) to the agent with the largest value for them; if this is not EF1, run the black-box mechanism.

The case of strictly positive valuations. It is interesting to notice that when valuations are strictly positive, i.e. $v_{i,j} > 0$ for all i, j , Mechanism 1 simply outputs the allocation of \mathcal{M}^* (Case IV). However, this does *not* imply that every PO and EF1 algorithm is NOM: in calculating worst-case outcomes, agents consider the possibility that others have zero valuations. This fact is critical in our proof of correctness (see the proof of Lemma 3).

Establishing the main result. Towards establishing Theorem 5, we begin by proving the following supporting lemmas. We show that Mechanism 1 preserves the efficiency and EF1 guarantees of \mathcal{M}^* .

Lemma 1. *If \mathcal{M}^* is α -Pareto efficient (resp. α -fractionally Pareto efficient) then Mechanism 1 always outputs an α -Pareto efficient (resp. α -fractionally Pareto efficient) partial allocation.*

Lemma 2. *If \mathcal{M}^* outputs non-wasteful, clean, and EF1 allocations then Mechanism 1 always outputs non-wasteful, clean, and EF1 partial allocations.*

The following lemma characterizes the set of allocations that Mechanism 1 could possibly return given the reported valuation of a particular agent. Specifically, we will show that every clean, non-wasteful and EF1 allocation, that is consistent with the reported valuation, is a possible output of Mechanism 1. Before stating the lemma, we define some useful notations. For any agent $i \in N$ and valuation function \mathbf{v}_i , let $\text{EF1}(i, \mathbf{v}_i)$ be the set of (partial) allocations $A = (A_1, \dots, A_n)$ that are clean, non-wasteful and envy-free up to one item with respect to agent i when her valuation function is \mathbf{v}_i , i.e., $\text{EF1}(i, \mathbf{v}_i) = \{A \in \Pi_n(\mathcal{M}) \mid \forall g \in A_i \ v_{i,g} > 0, \text{ and } \forall j \in N \text{ with } A_j \neq \emptyset, \mathbf{v}_i(A_i) \geq \mathbf{v}_i(A_j \setminus \{g\}) \text{ for some } g \in A_j, \text{ and } \mathbf{v}_i(\mathcal{M} \setminus \cup_{k \in N} A_k) = 0\}$.

Lemma 3. *Given any agent $i \in N$ and valuation function \mathbf{v}_i , for every allocation $A \in \text{EF1}(i, \mathbf{v}_i)$ there exists a set of valuations \mathbf{v}_{-i} such that Mechanism 1 on input $(\mathbf{v}_i, \mathbf{v}_{-i})$ outputs allocation A .*

Lemma 3 establishes that all partial allocations in $\text{EF1}(i, \mathbf{v}_i)$ can possibly be returned by the mechanism if agent i reports \mathbf{v}_i . Therefore, the problem of whether Mechanism 1 is obviously manipulable (partially) reduces to whether some set $\text{EF1}(i, \mathbf{v}'_i)$ has a better worst-case outcome than $\text{EF1}(i, \mathbf{v}_i)$, with respect to the valuation vector \mathbf{v}_i . Before describing the subsequent lemma which develops this idea, we define some relevant notations. Given any set of (partial) allocations $S \subseteq \Pi_n(\mathcal{M})$, an agent $i \in N$ and a valuation vector \mathbf{v} define $\ell(i, \mathbf{v}, S) = \min_{A \in S} \mathbf{v}(A_i)$.

Lemma 4. For any valuations \mathbf{v}, \mathbf{v}' and any agent $i \in N$, $\ell(i, \mathbf{v}, \text{EF1}(i, \mathbf{v})) \geq \ell(i, \mathbf{v}, \text{EF1}(i, \mathbf{v}'))$.

Proof. Let $A^* \in \text{EF1}(i, \mathbf{v})$ be an allocation such $\mathbf{v}(A_i^*) = \ell(i, \mathbf{v}, \text{EF1}(i, \mathbf{v}))$, i.e. A^* is a worst-case outcome for agent i when her valuation is \mathbf{v} . We can assume, without loss of generality, that A^* is a complete allocation because if we assign the unallocated items, $M \setminus \cup_{k \in N} A_k^*$, to agents other than agent i arbitrarily then still the resultant allocation would continue to belong in $\text{EF1}(i, \mathbf{v})$. This follows from the fact that A is non-wasteful, $\mathbf{v}(M \setminus \cup_{k \in N} A_k^*) = 0$.

If $A^* \in \text{EF1}(i, \mathbf{v}')$, the lemma immediately follows, since the minimum $\mathbf{v}(B_i)$, over all allocations $B \in \text{EF1}(i, \mathbf{v}')$, is at most $\mathbf{v}(A_i)$. Otherwise, $A^* \notin \text{EF1}(i, \mathbf{v}')$. Since A^* is complete, and thereby non-wasteful, this is possible if either bundle A_i is not clean wrt valuation \mathbf{v}' or agent i envies some other agent even after removal of one good. Denote by A the allocation obtained by cleaning the bundle A_i^* wrt valuation \mathbf{v}' , all other bundles remain as it is, i.e., $A_j = A_j^*$ for all $j \neq i$. Now if allocation $A \in \text{EF1}(i, \mathbf{v}')$, the lemma follows, because $\mathbf{v}(A_i^*) \geq \mathbf{v}(A_i)$ and $\ell(i, \mathbf{v}, \text{EF1}(i, \mathbf{v}'))$ is at most $\mathbf{v}(A_i)$. In the rest of the proof we will handle the case when $A \notin \text{EF1}(i, \mathbf{v}')$.

Towards this we will construct an allocation A' with the following properties: (i) $\mathbf{v}(A'_i) \leq \mathbf{v}(A_i)$, and (ii) $A' \in \text{EF1}(i, \mathbf{v}')$. These two properties together establish the desired inequality, since $\ell(i, \mathbf{v}, \text{EF1}(i, \mathbf{v})) = \mathbf{v}(A_i^*) \geq \mathbf{v}(A_i) \geq^{(i)} \mathbf{v}(A'_i) \geq^{(ii)} \ell(i, \mathbf{v}, \text{EF1}(i, \mathbf{v}'))$.

Let $S^* := \{j \in N \setminus \{i\} \mid A_j \neq \emptyset\}$ be the subset of agents with non-empty bundles in A . Define agent $j^* := \arg \max_{k \in S^*} \min_{g \in A_k} \mathbf{v}'(A_k \setminus \{g\})$, to be the agent in S^* whose allocation ‘‘up to one item’’ has maximum value with respect to \mathbf{v}' , and let $g^* = \arg \min_{g \in A_{j^*}} \mathbf{v}'(A_{j^*} \setminus \{g\})$ be the favourite good with respect to \mathbf{v}' that j^* has. A' is defined as follows. $A'_i = A_{j^*} \setminus \{g^*\}$, $A'_{j^*} = A_i \cup \{g^*\}$, and for all agents $k \in N \setminus \{i, j^*\}$ we have $A'_k = A_k$. It remains to show that A' satisfies (i) and (ii).

Towards proving (i) we have $\mathbf{v}(A'_i) = \mathbf{v}(A_{j^*} \setminus \{g^*\}) \leq \mathbf{v}(A_i)$, where the last inequality follows from the fact that $A \in \text{EF1}(i, \mathbf{v})$. For (ii), note that $N \setminus \{i\}$ can be written as the union of three disjoint sets S_1, S_2 and S_3 , where $S_1 = N \setminus (S^* \cup \{i\})$, $S_2 = (S^* \setminus \{j^*\})$, and $S_3 = \{j^*\}$. We will show that in allocation A' agent i with valuation vector \mathbf{v}' , does not envy, up to one item, any of the agents in S_1, S_2 , or S_3 ; this establishes that $A' \in \text{EF1}(i, \mathbf{v}')$. First, for every agent $k \in S_1$ we have that $A'_k = A_k = \emptyset$, by the definition of S^* , so agent i cannot envy any such agent. Second, given the definition of agent j^* , we know that for each agent k in $S_2 = S^* \setminus \{j^*\}$, we have

$$\begin{aligned} \min_{g \in A'_k} \mathbf{v}'(A'_k \setminus \{g\}) &= \min_{g \in A_k} \mathbf{v}'(A_k \setminus \{g\}) && (A'_k = A_k \text{ for all agents except } i \text{ and } j^*) \\ &\leq \min_{g \in A_{j^*}} \mathbf{v}'(A_{j^*} \setminus \{g\}) && \text{(by the definition of } j^*) \\ &\leq \mathbf{v}'(A_{j^*} \setminus \{g^*\}) && (g^* \in A_{j^*}) \\ &= \mathbf{v}'(A'_i). && (A'_i = A_{j^*} \setminus \{g^*\}) \end{aligned}$$

Therefore, agent i with valuation \mathbf{v}' does not envy up to one item (and, in fact, up to any item) any agent in S_2 . Third, recall that $A \notin \text{EF1}(i, \mathbf{v}')$. By definition, agent j^* has the bundle that \mathbf{v}' likes the most ‘‘up to one item’’, and g^* is that most valuable (in terms of marginal value) item in j^* ’s bundle, so $\mathbf{v}'(A_i) < \mathbf{v}'(A_{j^*} \setminus \{g^*\})$. Therefore, we have

$$\begin{aligned} \mathbf{v}'(A'_i) &= \mathbf{v}'(A_{j^*} \setminus \{g^*\}) && (A'_i = A_{j^*} \setminus \{g^*\}) \\ &> \mathbf{v}'(A_i) && (A \notin \text{EF1}(i, \mathbf{v}')) \\ &= \mathbf{v}'(A_i \cup \{g^*\} \setminus \{g^*\}) \\ &= \mathbf{v}'(A'_{j^*} \setminus \{g^*\}). && (A'_{j^*} = A_i \cup \{g^*\}) \end{aligned}$$

That is, agent i with valuation \mathbf{v}' does not envy, up to one item, any agent in $S_3 = \{j^*\}$. Note that, allocation A' is non-wasteful wrt agent i . Furthermore, the bundle A'_i can be cleaned maintaining the fact that A' is EF1 wrt agent i . This establishes that $A' \in \text{EF1}(i, \mathbf{v}')$ and concludes the proof. \square

Given the previous lemmas, we can show our main theorem with respect to NOM.

Theorem 6. If agents’ valuation functions are additive and M^* outputs EF1, non-wasteful and clean allocations, then Mechanism 1 is not obviously manipulable.

Proof. We first show that agents cannot improve their worst-case utilities by misreporting.

Let $i \in N$ be an agent whose true valuation is \mathbf{v}_i and let her misreport be \mathbf{b}_i . By Lemma 3, for every reported valuation vector \mathbf{b} of agent i , and for every allocation $A \in \text{EF1}(i, \mathbf{b})$, there exists a valuation profile \mathbf{v}_{-i} such that A is the output of Mechanism 1. On the other hand, by Lemma 2, Mechanism 1 always outputs allocations A that are EF1, clean, and non-wasteful. Therefore, every output allocation of Mechanism 1 is in $\text{EF1}(i, \mathbf{b})$. Therefore, if agent i reports \mathbf{v}_i (respectively \mathbf{b}_i), then the set of all possible allocations of Mechanism 1 is exactly $\text{EF1}(i, \mathbf{v}_i)$ (respectively $\text{EF1}(i, \mathbf{b}_i)$). Overloading notation, let $A_i^*(\mathbf{v}_i, \mathbf{v}_{-i})$ be the allocation of agent i in Mechanism 1, on input $(\mathbf{v}_i, \mathbf{v}_{-i})$. For the worst-case utility of agent i we have $\min_{\mathbf{v}_{-i}} \mathbf{v}_i(A_i^*(\mathbf{v}_i, \mathbf{v}_{-i})) = \min_{A \in \text{EF1}(i, \mathbf{v}_i)} \mathbf{v}_i(A_i) = \ell(i, \mathbf{v}_i, \text{EF1}(i, \mathbf{v}_i)) \geq^{Lem. 4} \ell(i, \mathbf{v}_i, \text{EF1}(i, \mathbf{b}_i)) = \min_{\mathbf{v}_{-i}} \mathbf{v}_i(A_i^*(\mathbf{b}_i, \mathbf{v}_{-i}))$.

This establishes Inequality (1) from the definition of NOM. To conclude the proof, it remains to show that the best-case utility of an agent is also not improved by misreporting. Note that the best case for agent $i \in N$ occurs when all other agents have no value for the items, i.e. $D_j = \emptyset$ for all $j \in N \setminus \{i\}$. Hence the sets $\{D_k\}_{k \in [n]}$ are pairwise-disjoint, and Mechanism 1 (via Case I) allocates the entirety of D_i to agent i , which is impossible to improve upon. \square

The proof of the above theorem along with the supporting lemmas establishes Theorem 5. We note that that it is not possible to improve upon this result by adding ex-ante fairness guarantees; see Appendix D.3 for details.

Acknowledgements

The authors are supported in part by an NSF CAREER award CCF-2144208, a Google Research Scholar Award, and by the Algorand Centres of Excellence program managed by Algorand Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of Algorand Foundation.

References

- [ABCM17] Georgios Amanatidis, Georgios Birmpas, George Christodoulou, and Evangelos Markakis. Truthful allocation mechanisms without payments: Characterization and implications on fairness. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 545–562, 2017.
- [ACGH20] Rediet Abebe, Richard Cole, Vasilis Gkatzelis, and Jason D Hartline. A truthful cardinal mechanism for one-sided matching. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2096–2113. SIAM, 2020.
- [AFRC⁺16] Haris Aziz, Aris Filos-Ratsikas, Jiashu Chen, Simon Mackenzie, and Nicholas Mattei. Egalitarianism of random assignment mechanisms. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 1267–1268, 2016.
- [AG18] Itai Ashlagi and Yannai A Gonczarowski. Stable matching mechanisms are not obviously strategy-proof. *Journal of Economic Theory*, 177:405–425, 2018.
- [AL21] Haris Aziz and Alexander Lam. Obvious manipulability of voting rules. In *International Conference on Algorithmic Decision Theory*, pages 179–193. Springer, 2021.
- [AMN20] R Pablo Arribillaga, Jordi Massó, and Alejandro Neme. On obvious strategy-proofness and single-peakedness. *Journal of Economic Theory*, 186:104992, 2020.
- [Azi20a] Haris Aziz. Justifications of welfare guarantees under normalized utilities. *ACM SIGecom Exchanges*, 17(2):71–75, 2020.
- [Azi20b] Haris Aziz. Simultaneously achieving ex-ante and ex-post fairness. In *International Conference on Web and Internet Economics*, pages 341–355. Springer, 2020.

- [BCIZ21] Nawal Benabbou, Mithun Chakraborty, Ayumi Igarashi, and Yair Zick. Finding fair and efficient allocations for matroid rank valuations. *ACM Transactions on Economics and Computation*, 9(4):1–41, 2021.
- [BEF21] Moshe Babaioff, Tomer Ezra, and Uriel Feige. Fair and truthful mechanisms for dichotomous valuations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5119–5126, 2021.
- [Bud11] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [BV21] Siddharth Barman and Paritosh Verma. Truthful and fair mechanisms for matroid-rank valuations. *arXiv preprint arXiv:2109.05810*, 2021.
- [CGG13] Richard Cole, Vasilis Gkatzelis, and Gagan Goel. Mechanism design for fair division: allocating divisible items without payments. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 251–268, 2013.
- [CKM⁺19] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum nash welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–32, 2019.
- [EK03] Lars Ehlers and Bettina Klaus. Coalitional strategy-proof and resource-monotonic solutions for multiple assignment problems. *Social Choice and Welfare*, 21(2):265–280, 2003.
- [FGP14] Eric Friedman, Ali Ghodsi, and Christos-Alexandros Psomas. Strategyproof allocation of discrete jobs on multiple machines. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 529–546, 2014.
- [FGPS19] Eric J Friedman, Vasilis Gkatzelis, Christos-Alexandros Psomas, and Scott Shenker. Fair and efficient memory sharing: Confronting free riders. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1965–1972, 2019.
- [FSV20] Rupert Freeman, Nisarg Shah, and Rohit Vaish. Best of both worlds: Ex-ante and ex-post fairness in resource allocation. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 21–22, 2020.
- [FTTZ16] Dimitris Fotakis, Dimitris Tsipras, Christos Tzamos, and Emmanouil Zampetakis. Efficient money burning in general domains. *Theory of Computing Systems*, 59(4):619–640, 2016.
- [FV19] Diodato Ferraioli and Carmine Ventre. Obvious strategyproofness, bounded rationality and approximation. In *International Symposium on Algorithmic Game Theory*, pages 77–91. Springer, 2019.
- [GM21] Jugal Garg and Aniket Murhekar. On fair and efficient allocations of indivisible goods. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.
- [GP15] Jonathan Goldman and Ariel D Procaccia. Spliddit: Unleashing fair division algorithms. *ACM SIGecom Exchanges*, 13(2):41–46, 2015.
- [GZH⁺11] Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. Dominant resource fairness: Fair allocation of multiple resource types. In *Nsdi*, volume 11, pages 24–24, 2011.
- [HPPS20] Daniel Halpern, Ariel D Procaccia, Alexandros Psomas, and Nisarg Shah. Fair division with binary valuations: One rule to rule them all. In *International Conference on Web and Internet Economics*, pages 370–383. Springer, 2020.
- [HR08] Jason D Hartline and Tim Roughgarden. Optimal mechanism design and money burning. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 75–84, 2008.

- [ILWM17] Nicole Immorlica, Brendan Lucier, Glen Weyl, and Joshua Mollner. Approximate efficiency in matching markets. In *International Conference on Web and Internet Economics*, pages 252–265. Springer, 2017.
- [KM02] Bettina Klaus and Eiichi Miyagawa. Strategy-proofness, solidarity, and consistency for multiple assignment problems. *International Journal of Game Theory*, 30(3):421–435, 2002.
- [KV19] Maria Kyropoulou and Carmine Ventre. Obviously strategyproof mechanisms without money for scheduling. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1574–1581, 2019.
- [Li17] Shengwu Li. Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–87, 2017.
- [OSH19] Josué Ortega and Erel Segal-Halevi. Obvious manipulations in cake-cutting. *arXiv preprint arXiv:1908.02988*, 2019.
- [PPS15] David C Parkes, Ariel D Procaccia, and Nisarg Shah. Beyond dominant resource fairness: Extensions, limitations, and indivisibilities. *ACM Transactions on Economics and Computation (TEAC)*, 3(1):1–22, 2015.
- [RF90] Günther Ruhe and Bernd Fruhwirth. ϵ -optimality for bicriteria programs and its application to minimum cost flows. *Computing*, 44(1):21–34, 1990.
- [Rot82] Alvin E Roth. The economics of matching: Stability and incentives. *Mathematics of operations research*, 7(4):617–628, 1982.
- [Sch96] James Schummer. Strategy-proofness versus efficiency on restricted domains of exchange economies. *Social Choice and Welfare*, 14(1):47–56, 1996.
- [TM20] Peter Troyan and Thayer Morrill. Obvious manipulations. *Journal of Economic Theory*, 185:104970, 2020.
- [ZP20] David Zeng and Alexandros Psomas. Fairness-efficiency tradeoffs in dynamic fair division. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 911–912, 2020.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] See “Our contribution” in intro.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A]
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Section 2.
 - (b) Did you include complete proofs of all theoretical results? [Yes] All our theorems and lemmas have proofs, either in the first 9 pages, or in the supplementary material.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]

- (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
- 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
- 5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Missing preliminaries

Allocations. A randomized allocation $\mathbf{R} = \{(p^z, A^z)\}_{z=1}^k$ is a probability distribution (or a lottery) over a set of integral allocations, i.e., for every $z \in [k]$, A^z is an integral allocation that occurs with a probability p^z . The sum of the probabilities is equal to one, $\sum_{z=1}^k p^z = 1$. The integral allocations A^1, A^2, \dots, A^k constitute the *support* of \mathbf{R} . For each randomized allocation $\mathbf{R} = \{(p^z, A^z)\}_{z=1}^k$ there exists a corresponding *expected fractional allocation* $X = \sum_{z=1}^k p^z \cdot A^z$, where for each agent $i \in \mathbb{N}$, the fraction of items given to agent i is $X_i = \sum_{z=1}^k p^z \cdot A_i^z$.⁷ Here $X_{i,j}$ can be thought of as the probability with which agent i is allocated item j in an integral allocation that is sampled from \mathbf{R} . For notational clarity, we will use the letters X, Y to denote allocations that are fractional or integral, A, B to denote allocations that are exclusively integral and \mathbf{R} for randomized allocations.

Mechanisms. A mechanism \mathcal{R} is *randomized* if for every reported valuations \mathbf{b} it outputs a randomized allocation, i.e. it returns integral allocations that are drawn from a probability distribution corresponding to a randomized allocation. Since every randomized allocation has an associated expected fractional allocation, the output of a randomized mechanism for reported valuations \mathbf{b} can also be interpreted as representing a fractional allocation. We use $\mathcal{R}(\mathbf{b}) = (\mathcal{R}_1(\mathbf{b}), \dots, \mathcal{R}_n(\mathbf{b}))$ to denote the *expected fractional allocation* that a randomized mechanism outputs given bids \mathbf{b} ; the vector $\mathcal{R}_i(\mathbf{b}) \in [0, 1]^m$ represents the probabilities with which agent i receives each item in the sampled integral allocation. For a randomized mechanism, $\mathbf{v}_i(\mathcal{R}_i(\mathbf{b}))$ denotes the expected utility of agent i in the sampled integral allocation when the input to the mechanism \mathcal{R} is \mathbf{b} .

To define NOM for a randomized mechanism \mathcal{R} , we compare the expected utilities $\mathbf{v}_i(\mathcal{R}_i(\cdot, \cdot))$ in inequality (1) and (2), instead of $\mathbf{v}_i(\mathcal{M}_i(\cdot, \cdot))$ as in the case of deterministic mechanisms. Notice that for randomized mechanisms, the definition of NOM takes an expectation over the randomness of the mechanism, and minimum/maximum are over the reports of other agents; we sometimes write “NOM in expectation” when referring specifically to a randomized mechanism.

B Missing from Section 3

The PS-Lottery algorithm is based on the well-known probabilistic serial algorithm, which outputs fractional allocations that are envy-free. On a high level, the PS-Lottery algorithm uses Birkhoff’s algorithm⁸ to implement the fractional allocation output by probabilistic serial as a randomized allocation (a lottery) over a set of EF1 allocations. For the sake of completeness, the PS-Lottery algorithm is formally described in Appendix B.1.

We begin by proving a lemma that highlights a connection between not obvious manipulability and randomized mechanisms that output ex-ante proportional allocations. We note that a similar observation is made in [OSH19] in the context of cake cutting.

Lemma 5. *Inequality (1) (the worst-case guarantee) is satisfied for every randomized mechanism \mathcal{R} that outputs ex-ante proportional allocations.*

Proof. Let \mathcal{R} be a randomized mechanism that outputs ex-ante proportional allocations. Consider an agent $i \in \mathbb{N}$ with true valuation \mathbf{v}_i .

Suppose agent i reports her true valuation \mathbf{v}_i . Since the mechanism outputs ex-ante proportional allocations, for every possible reports of other agents, \mathbf{b}_{-i} , the expected fractional allocation output by \mathcal{R} , $Y = (Y_1, Y_2, \dots, Y_n)$ will be such that $\mathbf{v}_i(Y_i) \geq \frac{1}{n} \mathbf{v}_i(M)$. As a consequence, when agent i reports \mathbf{v}_i , her worst-case expected utility is at least $\frac{1}{n} \mathbf{v}_i(M)$.

Next, we show that when every agent $j \in \mathbb{N} \setminus \{i\}$ reports her valuation to be \mathbf{v}_j (the true valuation of agent j), then the worst-case expected utility of agent i , as per her true valuation and irrespective of her report, will be at most $\frac{1}{n} \mathbf{v}_i(M)$, which implies the lemma. Consider the case when every agent $j \in \mathbb{N} \setminus \{i\}$ reports her valuation to be $\mathbf{b}_j = \mathbf{v}_j$ and agent i reports valuation \mathbf{b}_i . We know that the expected fractional allocation $Y' = (Y'_1, Y'_2, \dots, Y'_n)$ returned by \mathcal{R} will be proportional, i.e., for

⁷Note that multiple randomized allocations may have the same expected fractional allocation.

⁸Recall that, Birkhoff’s algorithm, given a square bistochastic matrix, decomposes it into a convex combination (or a lottery) over permutation matrices.

each agent $j \in N \setminus \{i\}$ we have $\mathbf{b}_j(Y'_j) = \mathbf{v}_i(Y'_j) \geq \frac{1}{n} \mathbf{v}_i(M)$. Adding up over all $j \in N \setminus \{i\}$ we have that $\sum_{j \in N \setminus \{i\}} \mathbf{v}_i(Y'_j) = \mathbf{v}_i(M) - \mathbf{v}_i(Y'_i) \geq \frac{n-1}{n} \mathbf{v}_i(M)$. On rearranging we get the required inequality $\mathbf{v}_i(Y'_i) \leq \frac{1}{n} \mathbf{v}_i(M)$. \square

Using Lemma 5 we can prove the following theorem.

Theorem 7. *The PS-Lottery algorithm of [Azi20b], which is ex-ante envy-free and ex-post EF1, is not obviously manipulable in expectation.*

Proof. Inequality (1) is implied by Lemma 5, since the PS-Lottery algorithm is ex-ante envy-free (see Appendix B.1), and therefore ex-ante proportional. It remains to prove Inequality (2).

First, it holds that the expected fractional allocation returned by the PS-Lottery algorithm, irrespective of the agents' reports, is such that each agent receives $\frac{m}{n}$ (fractional) items in total; see Property 2 in Appendix B.1. Second, the best-case for an agent $i \in N$ who, without loss of generality, values items in the order $v_{i,1} \geq v_{i,2} \dots \geq v_{i,m}$, and reports honestly, occurs when the reported valuation of other agents induce an opposite preference order on items, i.e., for each agent $j \in N \setminus \{i\}$ we have $b_{j,m} \geq b_{j,m-1} \dots \geq b_{j,1}$. In this case, agent i would receive items $1, 2, \dots, \lfloor \frac{m}{n} \rfloor$ in their entirety and a $\frac{m}{n} - \lfloor \frac{m}{n} \rfloor$ fraction of item $\lfloor \frac{m}{n} \rfloor + 1$. This allocation results in the maximum possible expected utility that agent i can get subject to the constraint that she gets $\frac{m}{n}$ fraction of items, and therefore it cannot be improved upon, no matter what her report is. \square

B.1 The PS-Lottery Algorithm of [Aziz, 2020b]

A square matrix $M \in [0, 1]^{k \times k}$ is *bistochastic* iff the sum of entries in each of its rows and columns is equal to one, i.e., for each $i \in [k]$, we have $\sum_{j=1}^k M_{i,j} = \sum_{j=1}^k M_{j,i} = 1$. Additionally, a bistochastic matrix $N \in \{0, 1\}^{k \times k}$ is a *permutation matrix* — each row and column of a permutation matrix contains exactly one entry having value one and all other entries are zero.

In essence, the PS-Lottery algorithm of [Azi20b] is based on the following two well-known algorithms:

Birkhoff's algorithm. Given a bistochastic matrix $M \in [0, 1]^{k \times k}$ as input, Birkhoff's algorithm can be used to decompose, in polynomial time, the matrix M into a convex combination of permutation matrices. That is, Birkhoff's algorithm outputs permutation matrices $\{M_i\}_{i=1}^t$ and positive real numbers $\{p_i\}_{i=1}^t$ such that $M = \sum_{i=1}^t p_i M_i$ and $\sum_i p_i = 1$; here $t = \mathcal{O}(k^2)$ is a positive integer.

Probabilistic serial algorithm. Given a fair division instance wherein agents have additive valuations, the probabilistic algorithm outputs a fractional allocation that is envy-free (and hence proportional). The fractional allocation output by probabilistic serial can be interpreted as the output of the following continuous procedure: starting from time $t = 0$, simultaneously, each agent start consuming items in the order of their preference (i.e., if $v_{i,j_1} \geq v_{i,j_2} \dots \geq v_{i,j_m}$ then the items are consumed in the order j_1, j_2, \dots, j_m ⁹) and at a rate of one item per unit time. If an item is fully consumed, then agents start consuming the next item as per their preference order. The algorithm terminates when all items have been consumed, which happens at time $t = \frac{m}{n}$. As a direct consequence of this, (i) each agent gets exactly $\frac{m}{n}$ fraction of items at the end, and (ii) the resultant fractional allocation is envy-free, since at every point agents are consuming their most-valued remaining item.

The PS-Lottery algorithm uses Birkhoff's algorithm to decompose the fractional allocation output by probabilistic serial into a convex combination over integral allocations, i.e., a randomized allocation. In addition, each integral allocation in the support of the randomized allocation is EF1.

In Section 3, towards showing that the PS-Lottery algorithm is NOM in expectation, we use the following properties of the PS-Lottery algorithm.

Property 1. *The randomized allocation output by the PS-Lottery algorithm is ex-ante EF and ex-post EF1.*

Property 2. *The total fraction of items that each agent gets in the expected fractional allocation returned by the PS-Lottery algorithm is $\frac{m}{n}$*

⁹ties can be broken arbitrarily

Additionally, suppose that agent $i \in N$ prefers the items in the order $v_{i,1} \geq v_{i,j_2} \dots \geq v_{i,j_m}$, and the preference order to all the other agents is opposite, i.e., for all agents $k \in [n] \setminus \{i\}$, we have $v_{k,m} \geq v_{k,m-1} \dots \geq v_{k,1}$, then in the fractional allocation output by probabilistic serial, agent i gets the items $1, 2, \dots, \lfloor \frac{m}{n} \rfloor$ entirely, and a fraction $\frac{m}{n} - \lfloor \frac{m}{n} \rfloor$ of the item $\lfloor \frac{m}{n} \rfloor + 1$.

C Missing from Section 4

Proof of Theorem 2. We prove the theorem for the case of two items; the proof can be easily generalized to hold for any number of items. Let \mathcal{M} be the utilitarian social welfare maximizing algorithm, coupled with any tie-breaking rule. Assume that both agents report a value of 1 for the first item, and zero for the other item, i.e. $b_{i,1} = 1$ and $b_{i,2} = 0$ for $i \in \{1, 2\}$. Given these reports, there must be an agent who gets item 1 with probability at least $\frac{1}{2}$ (if \mathcal{M} is deterministic then this probability will be exactly 1); assume that this is agent 1, without loss of generality.

Now, consider the case when the true valuation of agent 1 is $\mathbf{v}_1 = (\frac{2}{3} + \epsilon, \frac{1}{3} - \epsilon)$, for some small $\epsilon > 0$. Additionally, suppose that agent 1 reports her true valuation (i.e., $\mathbf{b}_1 = \mathbf{v}_1$) and agent 2's reported value for item 1, $b_{2,1} > \frac{2}{3} + \epsilon$. In this case, the utilitarian social welfare maximizing allocation gives the first item to agent 2. Therefore, the worst-case utility of agent 1 when she reports her true valuation is at most $\frac{1}{3} - \epsilon$. Next, consider the dishonest report $\mathbf{b}_1 = (b_{1,1}, b_{1,2}) = (1, 0)$. Given this, if $b_{2,1} < 1$, then agent 1 gets the first item for a utility of $\frac{2}{3} + \epsilon$. Otherwise if $b_{2,1} = 1$, then agent 1 gets the first item with probability at least $1/2$ (this follows from our choice of agent 1), thus, her expected utility is at least $\frac{1}{2}(\frac{2}{3} + \epsilon) = \frac{1}{3} + \frac{\epsilon}{2}$. In either case, her utility is strictly larger than $\frac{1}{3} - \epsilon$, her worst-case utility under honest reporting. Therefore, \mathcal{M} is obviously manipulable. \square

Proof of Theorem 4. We prove the theorem for the case of $n = 3$ agents and $m = 4$ items; we describe how our arguments can be adjusted to work for $n > 3$ agents and any number of items at the end of this proof. Let \mathcal{M} be a (possibly randomized) mechanism that *always* outputs a Nash social welfare maximizing allocation.

Let the true valuation of agent 1 be $\mathbf{v}_1 = (3.9, 3, 2, 0.9)$. The subsequent proof has two parts: first, we will show that if agent 1 reports her true valuation, then the worst-case utility is exactly 2, and second, if agent 1 reports $\mathbf{b}_1 = (2, 2, 1, 1)$, then the worst-case would be strictly more than 2; the theorem follows.

Worst-case utility when reporting honestly: Consider the case where agent 1 reports her true valuation, i.e., the report $\mathbf{b}_1 = \mathbf{v}_1$. We know that mechanism \mathcal{M} must output Nash social welfare maximizing allocations, and such allocations are necessarily EF1 [CKM⁺19]; randomized \mathcal{M} will output ex-post EF1 allocations. Consequently, agent 1 must be allocated at least one item, since otherwise some other agent will receive at least two items, and agent 1 will envy that agent even upon the removal of any one item. Furthermore, if agent 1 is allocated only item 4 the overall allocation can't be EF1, since agent 1 will envy (even upon the removal of any item) the agent that gets two of the first three items ($v_{1,4}$ is smaller than all other values). Thus, in the worst-case, agent 1 will get a bundle whose value is at least 2, her value for item 3. Next, we show that her worst-case utility is exactly equal to 2.

Let the reported valuations of agent 2 be $\mathbf{b}_2 = (0, 1, 0, 0)$ and of agent 3 be $\mathbf{b}_3 = (2, 0, 0, 1)$. Given this, the Nash welfare maximizing allocation is unique, and this allocation is such that agent 1 gets only item 3, for a total utility of 2. To see that this allocation is unique, first notice that item 2 must go to agent 2 (otherwise agent 2's utility, and therefore Nash social welfare, will be zero). Additionally, item 3 must go to agent 1 since she is the only agent with a non-zero value for it. The remaining items (1 and 4), which are valued only by agents 1 and 3, must be allocated in a way that the Nash social welfare of the resultant allocation is maximized. A simple case analysis shows that the unique allocation that maximizes Nash social welfare gives both items to agent 3. A consequence of having a unique Nash social welfare maximizing allocation is that mechanism \mathcal{M} must output it, irrespective of whether \mathcal{M} is randomized or deterministic.

Worst-case utility when misreporting: Consider the case where agent 1 misreports her valuation as $\mathbf{b}_1 = (2, 2, 1, 1)$. Given this report, the mechanism \mathcal{M} must allocate at least one item to agent 1; otherwise its allocation would not be EF1 (ex-post EF1 for if \mathcal{M} is randomized). We will show that

agent 1's allocation cannot be only item 3 nor only item 4 (i.e., a bundle of value 1 with respect to the reported values), irrespective of the reports of agents 2 and agent 3. Specifically, agent 1 must be allocated either (i) at least two items, (ii) only item 1, or (iii) only item 2. In each of these cases, the value of the items allocated to agent 1 (as per her true valuation) will be strictly more than 2, which completes the proof.

We proceed to show that agent 1 cannot be allocated just item 3 or just item 4 in any Nash social welfare maximizing allocation. Towards a contradiction, assume that in a Nash social welfare maximizing allocation A , agent 1 is allocated either only item 3 or only item 4. Since agent 1 is allocated only one item, there must be some other agent who is allocated at least two items. We consider the following exhaustive cases based on the items allocated to this other agent.

Case I: There is an agent $j \in \{2, 3\}$ having both items 1 and 2. In this case, the allocation A is not EF1 — and hence not maximizing Nash social welfare — since agent 1 envies agent j (with respect to the reported valuation \mathbf{b}_1), even upon the removal of one item.

Case II: There is an agent $j \in \{2, 3\}$ who gets one item from the set of items $\{1, 2\}$ and one item from $\{3, 4\}$. Without loss of generality assume that agent j is allocated items 2 and 3, and agent 1 gets item 4. We will show that in this case A can never be a Nash social welfare maximizing allocation.

Since allocation A maximizes Nash social welfare, it must be that transferring item 2 from agent j to agent 1 does not increase the Nash social welfare, i.e., the following inequalities must hold:

$$\begin{aligned} b_{1,4} \cdot (b_{j,2} + b_{j,3}) &\geq b_{j,2}(b_{1,3} + b_{1,4}) \\ 1 \cdot (b_{j,2} + b_{j,3}) &\geq b_{j,2} \cdot (1 + 1) && \text{(substituting } b_{1,3} = b_{1,4} = 1) \\ b_{j,3} &\geq b_{j,2} && (1) \end{aligned}$$

Similarly, transferring item 3 from agent j to agent 1 must also not increase the Nash social welfare:

$$\begin{aligned} b_{1,4} \cdot (b_{j,2} + b_{j,3}) &\geq b_{j,3} \cdot (b_{1,2} + b_{1,4}) \\ 1 \cdot (b_{j,2} + b_{j,3}) &\geq b_{j,3} \cdot (2 + 1) && \text{(substituting } b_{1,2} = 2 \text{ and } b_{1,4} = 1) \\ b_{j,2} &\geq 2b_{j,3} && (2) \end{aligned}$$

Combining inequalities (1) and (2) we have $b_{j,3} \geq 2b_{j,3}$, which can be true only if $b_{j,3} = 0$. However, since A allocates item 3 to agent j , and $b_{j,3} = 0$, but $b_{1,3} > 0$, A is not Nash social welfare maximizing (in fact, not even Pareto efficient).

The same argument can be extended to instances having $n > 3$ agents and $m = n + 1$ items by considering the case wherein agent 1's true valuation, $\mathbf{v}_1 = (3.9, 3, 3, \dots, 3, 2, 0.9)$. Here, the misreported valuation which leads to an improvement in her (expected) worst-case utility is $\mathbf{b}_1 = (2, 2, 2, \dots, 2, 1, 1)$. \square

C.1 Egalitarian Social Welfare

In this section we will show that any mechanism — randomized or deterministic — that maximizes egalitarian social welfare is obviously manipulable. The *egalitarian social welfare* of an (integral or fractional) allocation X , denoted as $\text{ESW}(X)$, is defined as the minimum utility that any agent derives from allocation X , i.e., $\text{ESW}(X) = \min_{i \in N} v_i(x_i)$. An integral allocation is *egalitarian social welfare maximizing* iff (i) it maximizes, among the set of all integral allocations, the number of agents having positive utility and (ii) for any such maximal set of agents S , it maximizes the egalitarian social welfare, i.e., the minimum utility of agents in S .

Our result rules out the existence of NOM mechanisms that output leximin allocations,¹⁰ since leximin allocations are, by definition, utilitarian social welfare maximizing.

Similar to utilitarian social welfare, when discussing egalitarian social welfare, we will assume that the valuations of agents are normalized: for every agent, the combined value for the set of all items is 1; see [AFRC⁺16] for a thorough discussion.

¹⁰An allocation is leximin iff it maximizes the lowest utility, subject to that the second lowest utility and so on.

For the case of $n = 2$ agents, there are no NOM mechanisms that maximize egalitarian social welfare for that case; this follows from Theorem 10 in Appendix E. The following theorem establishes that this impossibility continues to hold for the case of $n \geq 3$ agents.

Theorem 8. *Every (randomized or deterministic) mechanism that always outputs an allocation that maximizes the egalitarian social welfare is obviously manipulable, even for $n = 3$ agents and $m = 4$ items.*

Proof of Theorem 8. Consider any (randomized or deterministic) mechanism \mathcal{M} that maximizes the minimum utility among agents. To show that \mathcal{M} is obviously manipulable, we consider an instance with $n = 3$ agents and $m = 4$ items. Let the true (normalized) valuation of agent 1 be $\mathbf{v}_1 = (0.3, 0.3, 0.3, 0.1)$.

First, we will show that if agent 1 reports honestly, then the worst-case case outcome is that she is allocated only item 4, i.e., her utility in the worst-case is 0.1. Note that since \mathcal{M} maximizes the number of agents with positive utility, and agent 1 has a positive value for all 4 items, it must allocate at least one item to agent 1; otherwise either agent 2 or 3 has more than one item, one of which could be transferred to agent 1 to make her utility positive. Now, consider the case where the reports of other agents are $\mathbf{b}_2 = (0, 0, 1, 0)$ and $\mathbf{b}_3 = (0.05, 0.05, 0.9, 0)$. Here, item 3 must be allocated to agent 2 because it is the only item that she desires. By a straightforward case analysis, the unique way in which we can allocate the remaining items (items 1, 2 and 4) to maximize the minimum utility is by allocating item 4 to agent 1 and items 1 and 2 to agent 3. Since this allocation is unique, it must be the output of \mathcal{M} , irrespective of whether it is deterministic or randomized. Therefore, the utility of agent 1 is 0.1 under honest reporting, in the worst case.

Now, consider the case where agent 1 reports $\mathbf{b}_1 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0)$. Since \mathcal{M} maximizes the number of agents having non-zero utility, and agent 1 positively values the first three items, she must get at least one of them: otherwise some other agent will have at least two of them, and transferring one of them to agent 1 would strictly increase the number of agents with non-zero utility. Hence, the worst-case utility of agent 1 when reporting \mathbf{b}_1 but her true values are \mathbf{v}_1 is 0.3, her true value for each of the first three items. Since the worst-case utility improved, \mathcal{M} is obviously manipulable.

The same construction can be extended to $n > 3$ agents and $m = n + 1$ items by considering the case wherein agent 1's true valuation $\mathbf{v}_1 = (\frac{1}{n} - \epsilon, \frac{1}{n} - \epsilon, \frac{1}{n} - \epsilon, \dots, \frac{1}{n} - \epsilon, n \cdot \epsilon)$ for any positive constant $\epsilon < \frac{1}{n(n+1)}$, and she misreports to $\mathbf{b}_1 = (\frac{1}{n}, \frac{1}{n}, \frac{1}{n}, \dots, \frac{1}{n}, 0)$ leading to an improvement in her worst-case utility. \square

D Proofs missing from Section 5

D.1 Missing proofs

Proof of Lemma 1. First, note that Line 24 (making the temporary assignment clean) does not affect the utility of any agent i for her own bundle. Therefore if the allocation before this step was α -PO (respectively α -fPO) then after Line 24 these efficiency guarantees will continue to hold. Considering all possible cases for Mechanism 1, if all agents $j \in \mathbb{N}$ receive the bundle D_j (as in case I), the allocation is fractionally Pareto efficient. If all agents j receive the bundle D_j , with the exception of a single agent i that receives \widehat{D}_i , then all agents except i have the maximum possible utility, i.e. their utility for all items M . Furthermore, it is impossible to improve the utility of i without decreasing the utility of some other agent j ; this follows from the definition of \widehat{D}_i and D_j . Therefore, this allocation is fractionally Pareto efficient as well. This allocation is considered in cases II and III (if the black-box algorithm \mathcal{M}^* is not called), and it is the output of Mechanism 1 if it is also EF1. In the remaining cases (if the aforementioned allocation is not EF1, or we are in case IV) Mechanism 1 returns the allocation computed by the black-box algorithm \mathcal{M}^* (modulo Line 24), which is an α -PO (respectively α -fPO) partial allocation by definition. \square

Proof of Lemma 2. Line 24 does not affect the utility of any agent for her own bundle, and since valuations are additive, removal of items could only decrease her utility for the bundle of a different agent. Therefore if the allocation before this step was EF1, it remains EF1 afterwards. This step also ensures that the allocation is clean. Next, considering all possible cases, we first have that if every agent $j \in \mathbb{N}$ receives the bundle D_j (as in case I), the allocation is envy-free, and therefore EF1. In

cases II and III, Mechanism 1 either checks whether an allocation is EF1 before outputting it, or calls \mathcal{M}^* . In case IV \mathcal{M}^* is called. Whenever \mathcal{M}^* is called, then the output allocation is EF1 by definition.

Finally, note that Mechanism 1 returns non-wasteful allocations: it is easy to confirm that whenever Mechanism 1 doesn't call \mathcal{M}^* its allocation is non-wasteful, while if \mathcal{M}^* is called, the allocation of \mathcal{M}^* the returned allocation is non-wasteful and clean, so it remains non-wasteful after Line 24.¹¹ \square

Proof of Lemma 3. Given the valuation vector \mathbf{v}_i and an allocation $A = (A_1, A_2, \dots, A_n) \in \text{EF1}(i, \mathbf{v}_i)$ we will construct valuations for the other agents, \mathbf{v}_{-i} , such that Mechanism 1 outputs A on input $(\mathbf{v}_i, \mathbf{v}_{-i})$. The valuation of each agent $j \in (N \setminus \{i\})$ will be such that $v_{j,g} = 0$ for all $g \notin A_j$, and $v_{j,g} > 0$ for each good $g \in A_j$ (the precise value of $v_{j,g}$ will depend on how D_i intersects with A_j). That is, for each agent $j \in (N \setminus \{i\})$, $D_j = A_j$. Therefore, by construction, the subsets $\{D_j\}_{j=1}^n \setminus \{D_i\}$ are pairwise disjoint, and therefore $R_i = 1$ (so, we are never in case IV of Mechanism 1). Also, by construction, Line 24 will not affect any bundle. In the rest of the proof, we consider the three exhaustive cases, based on how many sets from $A_{-i} = (A_1, \dots, A_{i-1}, A_{i+1}, \dots, A_n)$ the set D_i intersects.

The first case is when D_i does not intersect any bundle from A_{-i} . Since $v_{i,g} = 0$ for every unallocated good g (by the definition of $\text{EF1}(i, \mathbf{v}_i)$), $D_i \subseteq A_i$, and therefore $R_j = 1$ for all $j \in N$. In this case we can set $v_{j,g}$ to be an arbitrary value for all $g \in A_j$, e.g. $v_{j,g} = 1$, for all agents $j \neq i$. Therefore, Mechanism 1 considers *Case I*, and it allocates the set D_j to every agent j . By construction, for all $j \neq i$, $D_j = A_j$. Additionally, for agent i , it must be that the bundle $A_i = D_i$, since A_i is also clean (by Line 24). Therefore, the output of Mechanism 1 is precisely A .

The second case is that D_i intersects exactly one bundle from A_{-i} . Let $j^* \in (N \setminus \{i\})$ be the (unique) agent such that $A_{j^*} \cap D_i \neq \emptyset$. Additionally, assume that $i < j^*$; an almost identical argument works if $i > j^*$. For each agent $k \in N \setminus \{i, j^*\}$, we define $v_{k,g} = 1$ if good $g \in A_k$ and $v_{k,g} = 0$ otherwise. The valuations of agent j^* are as follows

$$v_{j^*,g} = \begin{cases} 2\mathbf{v}_i(D_i \cap D_{j^*}) & g \in A_{j^*} \cap D_i \\ 1 & g \in A_{j^*} \setminus D_i \\ 0 & \text{otherwise} \end{cases}$$

Note that, the set of desired goods $D_k = A_k$ for all agents $k \in (N \setminus \{i\})$, and $\widehat{D}_i \supseteq A_i$. Furthermore, since A is non-wasteful (by the definition of $\text{EF1}(i, \mathbf{v}_i)$), goods in $\widehat{D}_i \setminus A_i$ must have zero value for i . Therefore, $\mathbf{v}_i(\widehat{D}_i) = \mathbf{v}_i(A_i)$. The construction satisfies that the bundles $\{D_k\}_{k \in N \setminus \{i\}}$ are pairwise disjoint, and the bundle D_i only intersects $D_{j^*} = A_{j^*}$. Therefore, $R_i = R_{j^*} = 1$ and $R_k = 0$ for every other agent k . Therefore, Mechanism 1, given valuations $(\mathbf{v}_i, \mathbf{v}_{-i})$, considers *Case III* (Lines 11-21) to compute the final allocation. In *Case III*, the mechanism first checks whether $\mathbf{v}_i(D_i \cap D_{j^*}) < \mathbf{v}_j(D_i \cap D_{j^*})$ holds, which is true for our construction since $v_{j^*,g} = 2\mathbf{v}_i(D_i \cap D_{j^*})$ for all $g \in A_{j^*} \cap D_i = D_{j^*} \cap D_i$. Afterwards, the mechanism checks if the allocation $(D_1, \dots, \widehat{D}_i, \dots, D_{j^*}, \dots, D_n)$ is EF1 (Line 13). Since $\mathbf{v}_i(\widehat{D}_i) = \mathbf{v}_i(A_i)$, and $A \in \text{EF1}(i, \mathbf{v}_i)$, this allocation is indeed EF1 for agent i , and since all other agents are envy-free, Mechanism 1 sets it as the temporary assignment. After removing all zero valued items from \widehat{D}_i (Line 24) we are left with the bundle A_i , hence, the final allocation is exactly A .

The third and final case is that D_i intersects more than one bundle from A_{-i} . Here, the valuation vector of every agent $j \in N \setminus \{i\}$ is such that $v_{j,g} = 1$ for $g \in A_j$ and $v_{j,g} = 0$ otherwise. By construction, the set of desired goods $D_j = A_j$ for every agent $j \in N \setminus \{i\}$, $\widehat{D}_i \supseteq A_i$, and since A is non-wasteful, we have $\mathbf{v}_i(\widehat{D}_i) = \mathbf{v}_i(A_i)$. D_i intersecting more than one bundle from A_{-i} , is equivalent to D_i intersecting more than one subsets from $\{D_j\}_{j \neq i}$, and therefore $R_j = 0$ for all $j \in N \setminus \{i\}$. Thus, given the valuation profile $(\mathbf{v}_i, \mathbf{v}_{-i})$, Mechanism 1 considers *Case II* (Lines 6-10), where it first checks whether the allocation $(D_1, D_2, \dots, \widehat{D}_i, \dots, D_n)$ is EF1. This allocation has the same value as A for agent i (and all other agents have exactly the same allocation), so agent i has an envy of at most one item by definition, and all other agents have no envy at all. Therefore,

¹¹Note that if the allocation of \mathcal{M}^* was non-wasteful and not clean, then Line 24 could have altered this allocation into a wasteful one.

Mechanism 1 sets this as the temporary assignment. After removing all zero valued items from \widehat{D}_i (Line 24) the final allocation is exactly A . This concludes the proof of Lemma 3. \square

D.2 Pseudo-code for Mechanism 1

REDUCTION 1: Black-box reduction

Input: Reported valuation functions of agents \mathbf{b} . Black-box access to an algorithm \mathcal{M}^* .
Output: A partial integral allocation $A = (A_1, \dots, A_n)$

- 1: **Set** $D_i \leftarrow \{g \in M \mid b_{i,g} > 0\}$.
- 2: **Set** $\widehat{D}_i \leftarrow M \setminus \cup_{j \neq i} D_j$ for each $i \in N$.
- 3: **Set** $R_i \leftarrow 1$ if subsets $\{D_j\}_{j=1}^n \setminus \{D_i\}$ are pairwise disjoint; 0 otherwise.
- 4: **if** the subsets $\{D_j\}_{j=1}^n$ are pairwise disjoint **then**
- 5: $A^* \leftarrow (D_1, \dots, D_n)$ } Case I
- 6: **else if** $\exists i \in N$ such that $R_i = 1$ and $R_j = 0$ for every $j \in N \setminus \{i\}$ **then**
- 7: **if** the allocation $(D_1, D_2, \dots, \widehat{D}_i, \dots, D_n)$ is EF1 **then**
- 8: $A^* \leftarrow (D_1, D_2, \dots, \widehat{D}_i, \dots, D_n)$ } Case II
- 9: **else**
- 10: $A^* \leftarrow \mathcal{M}^*(\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n)$
- 11: **else if** $\exists i, j \in N$ such that $i < j$ and $R_i = R_j = 1$, and $R_k = 0$ for $k \in [n] \setminus \{i, j\}$ **then**
- 12: **if** $\mathbf{b}_i(D_i \cap D_j) < \mathbf{b}_j(D_i \cap D_j)$ **then**
- 13: **if** $(D_1, \dots, \widehat{D}_i, \dots, D_j, \dots, D_n)$ is EF1 **then**
- 14: $A^* \leftarrow (D_1, \dots, \widehat{D}_i, \dots, D_j, \dots, D_n)$
- 15: **else**
- 16: $A^* \leftarrow \mathcal{M}^*(\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n)$ } Case III
- 17: **else**
- 18: **if** $(D_1, \dots, D_i, \dots, \widehat{D}_j, \dots, D_n)$ is EF1 **then**
- 19: $A^* \leftarrow (D_1, \dots, D_i, \dots, \widehat{D}_j, \dots, D_n)$
- 20: **else**
- 21: $A^* \leftarrow \mathcal{M}^*(\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n)$ } Case IV
- 22: **else**
- 23: $A^* \leftarrow \mathcal{M}^*(\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n)$
- 24: For all $i \in N$, iteratively remove goods $g \in A_i^*$ such that $\mathbf{b}_i(A_i^* \setminus \{g\}) = \mathbf{b}_i(A_i^*)$.
- 25: **return** A^*

D.3 Ex-ante EF, Ex-post fPO, Ex-post EF1 and NOM

In the previous section we established (Application 1) the existence of a deterministic NOM mechanism that outputs ex-post fPO and EF1 allocations for additive agents. Here we make a short remark that it is not possible to improve upon this result by adding ex-ante fairness guarantees. Specifically, there does not exist a randomized mechanism that is NOM in expectation and outputs ex-post fPO and EF1 allocations, that is additionally ex-ante envy-free. This follows directly from the following impossibility result of [FSV20].

Theorem 9 ([FSV20]). *There exists instances with additive valuation where there is no randomized allocation that is simultaneously ex-post fPO and EF1, and ex-ante envy-free.*

E Best of Both Worlds

In the previous sections we saw how fairness guarantees can be turned into NOM guarantees. For example, Lemma 5 says that the worst-case guarantee of NOM is satisfied for ex-ante proportional algorithms, while Theorem 5 turns any EF1 algorithm into a NOM (plus EF1) mechanism. In this section we show that this connection can be exploited in the other direction to prove impossibility results for fair algorithms. Specifically, we show that certain “best-of-both-worlds” fairness guarantees,

that is, randomized allocations that satisfy an ex-ante and ex-post guarantees simultaneously, are not guaranteed to exist. Our impossibility result uses the following theorem.

Theorem 10. *Every deterministic (or randomized) mechanism for $n = 2$ agents with additive and normalized utilities, that always outputs allocations that (ex-post) maximize the number of agents having positive utility, is obviously manipulable.*

Proof of Theorem 10. Let \mathcal{M} be such a deterministic (or randomized) mechanism. We will prove the theorem (i.e., show that \mathcal{M} is obviously manipulable) for the case of $m = 2$ items; the proof easily generalizes to more than 2 items.

Consider the case when both the agents report a value of 1 for the first item and zero for the other item, i.e. $\mathbf{b}_i = (1, 0)$ for $i \in \{1, 2\}$. Given these reports, there must be an agent who gets item 1 with probability at least $\frac{1}{2}$ (if \mathcal{M} is deterministic then this probability will be exactly 1); assume that this is agent 1, without loss of generality.

Consider the case that the true valuation of agent 1 is $\mathbf{v}_1 = (\frac{2}{3} + \epsilon, \frac{1}{3} - \epsilon)$, for some small constant $\epsilon > 0$. The worst case for agent 1 under honest reporting happens when the reported valuation of agent 2 $\mathbf{b}_2 = (1, 0)$, in which case item 1 must be allocated to agent 2 and item 2 to agent 1 – irrespective of whether the mechanism \mathcal{M} is deterministic or randomized since \mathcal{M} maximizes the number of agents having positive utility *ex-post*, and allocating item 1 to agent 2 and item 2 to agent 1 is the unique allocation that gives both agents non-zero utility. Therefore, the worst-case utility of agent 1 is at most $\frac{1}{3} - \epsilon$.

Now, consider the dishonest report $\mathbf{b}_1 = (b_{1,1}, b_{1,2}) = (1, 0)$. If $b_{2,2} > 0$, then mechanism \mathcal{M} , in order to maximize the number of agents having positive utility, will allocate item 2 to agent 2 and item 1 to agent 1, i.e., agent 1’s utility in this case (per \mathbf{v}_1) would be $\frac{2}{3} + \epsilon$. On the other hand, if $b_{2,2} = 0$ (and therefore $b_{2,1} = 1$) then as per our choice of mechanism \mathcal{M} , agent 1 will receive item 1 with probability at least $\frac{1}{2}$. Hence, her utility in this case would be at least $\frac{1}{2}\mathbf{v}_{1,1} = \frac{1}{3} + \frac{\epsilon}{2}$. The above case analysis shows that if agent 1 reports valuation \mathbf{b}_1 then her worst-case utility will be $\frac{1}{3} + \frac{\epsilon}{2}$, which is strictly larger than her worst-case utility when reporting true valuation. Therefore, \mathcal{M} is obviously manipulable. \square

We note that we can apply Theorem 10 to show that for $n = 2$ there is no NOM mechanism that outputs (ex-post) and MNW solution, since MNW is scale free (and therefore normalization comes without loss of generality).

We are now ready to prove our main result for this section.

Theorem 11. *Under additive valuations, randomized allocations that are ex-ante proportional, ex-post Pareto efficient and ex-post maximize the number of agents having positive utility do not exist, for any number of agents and any number of items.*

Proof. We begin by establishing the result for $n = 2$ agents. Towards a contradiction, assume that for additive valuations, randomized allocations that are ex-ante proportional, ex-post Pareto efficient and ex-post maximize the number of agents having positive utility always exists, and let \mathcal{R} be a randomized mechanism that, given the reported valuations of agents, outputs such a allocation. Note that this mechanism need not be computationally efficient. We will show that \mathcal{R} is NOM in expectation, in direct contradiction to Theorem 10.

Since, \mathcal{R} outputs ex-ante proportional allocations, the worst-case guarantee of NOM (inequality (1)) holds, via Lemma 5. Additionally, the best-case guarantee (inequality (2)) is also satisfied: if an agent $i \in \mathbb{M}$ reports her true valuation, then in the best case, all other agents only desire her least-valued item; in this case, agent i receives all items except her least-valued item. This outcome cannot be (strictly) improved as agent i must lose at least one item: the mechanism maximizes the number of agents having positive utility and each agent reports having a positive utility of at least one item, therefore, all items cannot be allocated to the same agent. Therefore, \mathcal{R} is NOM in expectation, a contradiction.

Thus, for $n = 2$ agents randomized allocation that are ex-ante proportional, ex-post Pareto efficient, and ex-post maximize the number of agents having positive utility do not always exist. In particular, there is a 2 agent instance where such randomized allocations do not exist. This instance can be easily converted into a $n > 2$ agent instance by adding dummy agents a_1, a_2, \dots, a_k and dummy

items j_1, j_2, \dots, j_k such that (i) for each $i \in [k]$, agent a_i 's value for item j_1 is 1 and her value for every other item is 0, (ii) the existing agents have zero value for each dummy item. In this new instance, every Pareto efficient allocation must be such that agent a_i is allocated item j_i and no other item that existing agents value positively. Thus, a randomized allocation that is ex-ante proportional, ex-post Pareto efficient, and ex-post maximizes the number of agents having positive utility would continue to satisfy these properties after removing the dummy agents and dummy items. However, as already established, such allocations do not exist for $n = 2$ agents. This concludes the proof of Theorem 11. \square

The existence of ex-ante proportional, ex-post Pareto efficient and ex-post EF1 allocations remains an elusive open problem. Our result shows that replacing ex-post EF1 with a different, mild fairness guarantee, namely ex-post maximizing the number of agents with positive utility, is impossible. The following corollary is immediate (and we note that point #3 in the following corollary is shown by [FSV20]).

Corollary 2. *Under additive valuations, randomized allocations that satisfy the following properties do not always exist.*

1. *ex-ante proportional, ex-post Pareto efficient and ex-post egalitarian welfare maximizing*
2. *ex-ante proportional and ex-post leximin allocations*
3. *ex-ante proportional and ex-post MNW.*