The Unreasonable Fairness of Maximum Nash Welfare

IOANNIS CARAGIANNIS, University of Patras, Greece

DAVID KUOKAWA, Carnegie Mellon University, USA

HERVÉ MOULIN, University of Glasgow, UK and Higher School of Economics, St Petersburg, Russia

ARIEL D. PROCACCIA, Carnegie Mellon University, USA

NISARG SHAH, University of Toronto, Canada

JUNXING WANG, Carnegie Mellon University, USA

The *maximum Nash welfare* (MNW) solution—which selects an allocation that maximizes the product of utilities—is known to provide outstanding fairness guarantees when allocating divisible goods. And while it seems to lose its luster when applied to indivisible goods, we show that, in fact, the MNW solution is strikingly fair even in that setting. In particular, we prove that it selects allocations that are envy-free up to one good—a compelling notion that is quite elusive when coupled with economic efficiency. We also establish that the MNW solution provides a good approximation to another popular (yet possibly infeasible) fairness property, the maximin share guarantee, in theory and—even more so—in practice. While finding the MNW solution is computationally hard, we develop a nontrivial implementation and demonstrate that it scales well on real data. These results establish MNW as a compelling solution for allocating indivisible goods and underlie its deployment on a popular fair-division website.

CCS Concepts: • **Theory of computation** → **Algorithmic mechanism design**; • **Applied computing** → **Economics**;

Additional Key Words and Phrases: Fair division, resource allocation, Nash welfare

ACM Reference format:

Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Trans. Econ. Comput.* 7, 3, Article 12 (September 2019), 32 pages.

<https://doi.org/10.1145/3355902>

A preliminary version of this article appeared in proceedings of the 17th ACM Conference on Economics and Computation, pp. 305–322, 2016.

This work was supported in part by the National Science Foundation under grants IIS-1350598, CCF-1215883, and CCF-1525932; by a Sloan Research Fellowship; by the COST Action IC1205 on “Computational Social Choice”; by the Caratheodory research grant E.114 from the University of Patras; and by the NSERC under the Discovery Grants program. Authors’ addresses: I. Caragiannis, University of Patras, Panepistimioupoli Patron, Patras, 265 04, Greece; email: caragian@ceid.upatras.gr; D. Kurokawa, A. D. Procaccia, and J. Wang, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA, 15213; emails: {dkurokaw, arielpro, junxingw}@cs.cmu.edu; H. Moulin, University of Glasgow, University Avenue, Glasgow, G12 8QQ, UK, Higher School of Economics, St Petersburg, 16 Soyuza Pechatnikov St, St. Petersburg, 190121, Russia; email: herve.moulin@glasgow.ac.uk; N. Shah, University of Toronto, 10 King’s College Rd, Toronto, ON, M5G2R3, Canada; email: nisarg@cs.toronto.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2167-8375/2019/09-ART12 \$15.00

<https://doi.org/10.1145/3355902>

1 INTRODUCTION

We are interested in the problem of fairly allocating *indivisible* goods, such as jewelry or artworks. But to better understand the context for our work, let us start with an easier problem: fairly allocating *divisible* goods. Specifically, let there be m homogeneous divisible goods, and n players with *linear* valuations over these goods; that is, if player i receives an x_{ig} fraction of good g , her value is $v_i(\mathbf{x}_i) = \sum_g x_{ig} v_i(g)$, where $v_i(g)$ is her non-negative value for the (entire) good g alone.

The question, of course, is what fraction of each good to allocate to each player; and it has an elegant answer, given more than four decades ago by Varian [41]. Under his *competitive equilibrium from equal incomes (CEEI)* solution, all players are endowed with an equal budget, say \$1 each. The equilibrium is an allocation coupled with (virtual) prices for the goods, such that each player buys her favorite bundle of goods for the given prices, and the market clears (all goods are sold). One formal way to argue that this solution is fair is through the compelling notion of *envy freeness* [23]: Each player weakly prefers her own bundle to the bundle of any other player. This property is obviously satisfied by CEEI, as each player can afford the bundle of any other player, but instead chooses to buy her own bundle.

While the CEEI solution may seem technically unwieldy at first glance, it always exists, and, in fact, has a very simple structure in the foregoing setting: the CEEI allocations (which are what we care about, as the prices are virtual) exactly coincide with allocations $\mathbf{x}$ that maximize the *Nash social welfare* $\prod_i v_i(\mathbf{x}_i)$ [5, Volume 2, Chapter 14]. Consequently, a CEEI allocation can be computed in polynomial time via the convex program of Eisenberg and Gale [22].

Let us now revisit our original problem—that of allocating *indivisible* goods, under *additive valuations*: the utility of a player for her bundle of goods is simply the sum of her values for the individual goods she receives. This is an inhospitable world where central fairness notions like envy freeness cannot be guaranteed (just think of a single indivisible good and two players). Needless to say, the existence of a CEEI allocation is no longer assured.

Nevertheless, the idea of maximizing the Nash social welfare (that is, the product of utilities) seems natural in and of itself [19, 37]. Informally, it hits a sweet spot between Bentham’s utilitarian notion of social welfare—maximize the sum of utilities—and the egalitarian notion of Rawls—maximize the minimum utility. Moreover, this solution is *scale-free*, in the sense that scaling a player’s valuation function would not change the outcome [31]. But, when the maximum Nash welfare solution is wrenched from the world of divisible goods, it seems to lose its potency. Or does it?

Our goal in this article is to demonstrate the “unreasonable effectiveness” [43]—or unreasonable fairness, if you will—of the *maximum Nash welfare (MNW)* solution, even when the goods are indivisible. We wish to convince the reader that

... the MNW solution exhibits an elusive combination of fairness and efficiency properties, and can be easily computed in practice. It provides the most practicable approach to date—arguably, the ultimate solution—for the division of indivisible goods under additive valuations.

1.1 Real-world Connections and Implications

Our quest for fairer algorithms is part of the growing body of work on practical applications of computational fair division [1, 16, 24, 27, 36]. We are especially excited about making a real-world impact through *Spliddit* (www.spliddit.org), a not-for-profit fair-division website [25]. Since its launch in November 2014, the website has attracted more than 90K users. The motto of Spliddit is *provably fair solutions*, meaning that the solutions obtained from each of the website’s five applications satisfy guaranteed fairness properties. These properties are carefully explained to users,

thereby helping users understand why the solutions are fair and increasing the likelihood that they would be adopted (in contrast, trying to explain the algorithms themselves would be much trickier).

One of Spliddit's five applications is allocating goods. In our view, it is the hardest problem Spliddit attempts to solve, and the previous solution left something to be desired. Here is how it worked: First, to express their preferences, users simply need to divide 1K points between the goods. This simple elicitation process relies on the assumption of additive preferences, and is the reason why, in our view, this assumption is indispensable in practical applications. Given these inputs, the algorithm considers three levels of fairness: envy freeness (explained above), proportionality (each player receives $1/n$ of her value for all the goods), and *maximin share guarantee* (each player i receives a bundle worth at least $\max_{X_1, \dots, X_n} \min_j v_i(X_j)$, where $X_1, \dots, X_n$ is a partition of the goods into n bundles). The algorithm finds the highest feasible level of fairness, and subject to that, maximizes utilitarian social welfare. Importantly, a *maximin share allocation* (which gives each player her maximin share guarantee) may not exist, but a $(2/3)$ -approximation thereof is always feasible; that is, each player can receive at least $2/3$ of her maximin share guarantee [36]. This allowed Spliddit to provide a *provable* fairness guarantee for indivisible goods. That said, a (full) maximin share allocation can always be found in practice [12, 28].

While the algorithm generally provided good solutions, it was highly discontinuous, and its direct reliance on the maximin share alone—when envy freeness and proportionality cannot be obtained—sometimes led to nonintuitive outcomes. For example, consider this excerpt from an email sent by a Spliddit user on January 7, 2016:

"Hi! Great app :) We're 4 brothers that need to divide an inheritance of 30+ furniture items. This will save us a fist fight :) I played around with the demo app and it seems there are non-optimal results for at least two cases where everyone distributes the same amount of value onto the same goods.... Try 3 people, 5 goods, with everyone placing 200 on every good.... [This] case gives 3 to one person and 1 to each of the others. Why is that?"

The answer to the user's question is that envy freeness and proportionality are infeasible in the example, so the algorithm sought a maximin share allocation. In every partition of the five goods into three bundles there is a bundle with at most one good (worth 200 points), hence the maximin share guarantee of each player is 200 points. Therefore, giving three goods to one player and one good to each of the others indeed maximizes utilitarian social welfare subject to giving each player her maximin share guarantee. Note that the MNW solution produces the intuitively fair allocation in this example (two players receive two goods each, one player receives one good).

Based on the results described below, we firmly believe that the MNW solution is superior to the previous algorithm for allocating goods (and to every other approach we know of, as we discuss below). It has been deployed on Spliddit since May 24, 2016.

1.2 Our Results

In order to circumvent the possible nonexistence of envy-free allocations, we consider a slightly relaxed version, *envy freeness up to one good* (EF1) [30]. In an allocation satisfying this property, player i may envy player j , but the envy can be eliminated by removing a single good from the bundle of player j . We show that the MNW solution always outputs an allocation that is envy-free up to one good, as well as Pareto optimal—a well-known notion of economic efficiency. And while envy freeness up to one good is straightforward to obtain in isolation, achieving it together with Pareto optimality is challenging; the fact that the MNW solution does so is a strong argument in its favor. In particular, as discussed in Section 1.1, on Spliddit it is crucial to be able to explain to

users what the guarantees of each method are; in our view, these two properties are especially compelling and easy to understand.

As another measure for the fairness of the MNW solution, we study the *maximin share* property. As mentioned earlier, the algorithm currently deployed on Spliddit relies on the existence of an approximate version of this property [36]. With this in mind, we show that the MNW solution always guarantees each of the n players a π_n -fraction of her maximin share guarantee, where $\pi_n = 2/(1 + \sqrt{4n - 3})$. Strikingly, this ratio is completely tight. Furthermore, we introduce a novel and equally attractive variant, *pairwise maximin share*, which is incomparable to the original property. Using the previous result, we prove that under the MNW solution, each player receives at least a Φ -fraction of her pairwise maximin share guarantee, where $\Phi = (\sqrt{5} - 1)/2 \approx 0.618$ is the golden ratio conjugate, and that this ratio is also tight. Experiments provide further evidence in favor of the MNW solution: It gives an excellent approximation to both MMS and pairwise MMS in practice. Among the 1,281 real-world fair-division instances we recorded on Spliddit, it achieves full MMS and pairwise MMS on more than 95% and 90% of the instances, respectively, and never worse than a 3/4-approximation on any instance.

The problem of computing an MNW allocation is known to be strongly $\mathcal{NP}$ -hard [33]. One of our main contributions is the algorithm we devised for computing an MNW allocation for the form of valuations elicited on Spliddit, in which a player is required to divide 1K points among the available goods. Our algorithm scales very well, solving relatively large instances with 50 players and 150 goods in less than 30s, while other candidate algorithms we describe fail to solve even small instances with 5 players and 15 goods in twice as much time.

1.3 Related Work

The concept of envy freeness up to one good originates in the work of Lipton et al. [30]. They deal with general combinatorial valuations and give a polynomial-time algorithm that guarantees that the maximum envy is bounded by the maximum marginal value of any player for any good; this guarantee reduces to EF1 in the case of additive valuations. However, in the additive case, EF1 alone can be achieved by simply allocating the goods to players in a round-robin fashion, as we discuss below. The algorithm of Lipton et al. [30] does not guarantee additional properties.

Budish [16] introduces the concept of *approximate CEEI*, which is an adaptation of CEEI to the setting of indivisible goods (among other contributions in this beautiful paper, he also introduces the notion of maximin share guarantee). He shows that an approximate CEEI exists and (approximately) guarantees certain properties. The approximation error goes to zero when the number of goods is fixed, whereas the number of players, as well as the number of copies of each good, go to infinity. His approach is practicable in the MBA course allocation setting, which motivates his work—there are many students, many seats in each course, and relatively few courses. But it does not give useful guarantees for the type of instances we encounter on Spliddit, where the number of players is small, and there is typically one copy of each good.

Maximizing Nash welfare is known to be appealing in settings with divisible goods. Varian [41] shows that in a setting almost identical to ours, except that each good is perfectly divisible, maximizing Nash welfare produces a CEEI allocation, which implies that it is envy-free and Pareto optimal. Weller [42] generalizes this to the cake-cutting setting [39], in which a heterogeneous good is to be divided. For cake cutting, Berliant et al. [8] strengthen the fairness guarantee by showing that maximizing Nash welfare satisfies “group envy-freeness,” and Sziklai and Segal-Halevi [40] show that maximizing Nash welfare satisfies intuitive resource and population monotonicity properties.

Little is known about the axiomatic properties of maximizing Nash welfare with *indivisible* goods. From an algorithmic perspective, Ramezani and Endriss [37] show that maximizing Nash

welfare is $\mathcal{NP}$ -hard under certain combinatorial bidding languages (including, under additive valuations). Cole and Gkatzelis [19] give a polynomial-time constant-factor approximation under additive valuations (to be precise, their objective function is the geometric mean of the utilities). Anari et al. [4] and Cole et al. [18] improve the approximation ratio, and Anari et al. [3] provide a constant-factor approximation for more general piecewise-linear concave valuations. Lee [29] shows that the problem is APX-hard; that is, a constant-factor approximation is the best one can hope to achieve in polynomial time. However, a constant-factor approximation need not satisfy any of the theoretical guarantees we establish in this article for the MNW solution.

When there are only two players, other compelling approaches for allocating goods are available. In fact, Spliddit used to handle this case separately, via the Adjusted Winner algorithm [15]. The shortcoming of Adjusted Winner is that, except in knife-edge situations, it has to split one of the goods between the two players. Adjusted Winner can be interpreted as a special case of the Egalitarian Equivalent rule of Pazner and Schmeidler [35], which is defined for any number of players. For $n > 2$ players, it may need to split up to $n - 1$ goods (or all the goods if $m < n$); thus, it is impractical to apply it to indivisible goods.

Let us briefly mention two additional models for the division of indivisible goods. First, some papers assume that the players express *ordinal* preferences (i.e., a ranking) over the goods [6, 11, 13, 14]. This assumption (arguably) does not lead to crisp fairness guarantees—the goal is typically to design algorithms that compute fair allocations if they exist. Second, it is possible to allow randomized allocations [9, 10, 17]; this is hardly appropriate for the cases we find on Spliddit in which the outcome is used only once. We also remark that in work that builds on the conference version of this article, Conitzer et al. [20] study a public decision setting that is more general than our indivisible goods allocation setting, and establish attractive properties (though weaker than ours) of the MNW allocation.

Finally, it is worth noting that the idea of maximizing the product of utilities was studied by Nash [32], in the context of his classic bargaining problem. This is why this notion of social welfare is named after him. In the networking community, the same solution goes by the name of *proportional fairness*, due to another property that it satisfies when goods are divisible [26]: When switching to any other allocation, the total percentage gains for players whose utilities increased sum to at most the total percentage losses for players whose utilities decreased; thus, in some sense, no such switch would be socially preferable.

2 MODEL

Let $[k] \triangleq \{1, \dots, k\}$. Let $\mathcal{N} = [n]$ denote the set of *players*, and $\mathcal{M}$ denote the set of *goods* with $m = |\mathcal{M}|$. Throughout the article, we assume the goods to be indivisible (i.e., each good must be entirely allocated to a single player), but our method and its guarantees extend seamlessly to the case where some of the goods are divisible (see Section 6).

Each player i is endowed with a *valuation* function $v_i : 2^{\mathcal{M}} \rightarrow \mathbb{R}_{\geq 0}$ such that $v_i(\emptyset) = 0$. With the exception of Section 3.1, throughout the article, we assume that players' valuations are *additive*: $\forall S \subseteq \mathcal{M}, v_i(S) = \sum_{g \in S} v_i(\{g\})$. To simplify notation, we write $v_i(g)$ instead of $v_i(\{g\})$ for a good $g \in \mathcal{M}$. The assumption of additive valuations is common in the literature on the fair allocation of indivisible goods [12, 36]. Furthermore, eliciting more general combinatorial preferences is often difficult in practice, which is why, to our knowledge, all of the deployed implementations of fair-division methods for indivisible goods—including Adjusted Winner [15] and the algorithm implemented on Spliddit (see Section 1.1)—also rely on additive valuations. That said, our main result (Theorem 3.2) generalizes to more expressive submodular valuations (see Section 3.1).

Given the valuations of the players, we are interested in finding a *feasible allocation*. For a set of goods $S \subseteq \mathcal{M}$ and $k \in \mathbb{N}$, let $\Pi_k(S)$ denote the set of ordered partitions of S into k bundles. A

feasible allocation $\mathbf{A} = (A_1, \dots, A_n) \in \Pi_n(\mathcal{M})$ is a partition of the goods that assigns a subset A_i of goods to each player i . Under this allocation, the *utility* to player i is $v_i(A_i)$ (her value for the set of goods she receives).

Our goal is to find a *fair* allocation. The fair-division literature often takes an axiomatic approach to defining fairness; the most compelling definition is *envy freeness*.

Definition 2.1 (EF: Envy Freeness). An allocation $\mathbf{A} \in \Pi_n(\mathcal{M})$ is called envy-free if for all players $i, j \in \mathcal{N}$, we have $v_i(A_i) \geq v_i(A_j)$. That is, each player values her own bundle at least as much as she values any other player's bundle.

Envy freeness cannot be guaranteed in general; for example, allocating a single indivisible good among two players who value it positively would inevitably result in envy. In fact, it is computationally hard to determine whether an EF allocation exists [21]. To guarantee existence, a somewhat weaker definition is called for; the following definition is a rather minimal relaxation that is still interesting when the number of goods is larger than the number of players.

Definition 2.2 (EF1: Envy Freeness Up to One Good). An allocation $\mathbf{A} \in \Pi_n(\mathcal{M})$ is called envy-free up to one good (EF1) if¹

$$\forall i, j \in \mathcal{N}, \exists g \in A_j, v_i(A_i) \geq v_i(A_j \setminus \{g\}).$$

In words, i may envy j , but the envy can be eliminated by removing a single good from the bundle of j . More generally, one can define envy freeness up to k goods for every $k \in \mathbb{N}$, but as we show in this article, EF1 can always be guaranteed along with other desirable properties, eliminating the need to relax the requirement further.

Another relaxation of envy freeness is known as the *maximin share guarantee* [16]. It is a natural extension of the 2-player cut-and-choose idea to the case of n players. Informally, the maximin share guarantee of a player is the value she can secure if she were allowed to divide the set of goods into n bundles, but then chose a bundle last (thus, possibly ending up with her least-valued bundle).

Definition 2.3 (MMS: Maximin Share). The maximin share (MMS) guarantee of player i is given by

$$\text{MMS}_i = \max_{\mathbf{A} \in \Pi_n(\mathcal{M})} \min_{k \in [n]} v_i(A_k).$$

We say that $\mathbf{A}$ is an α -MMS allocation if $v_i(A_i) \geq \alpha \cdot \text{MMS}_i$ for all players $i \in \mathcal{N}$.

Note that, in principle, MMS_i depends on v_i and n ; these parameters are not part of the notation, as they will always be clear from the context. While it is impossible to guarantee all players their full maximin share [28, 36], a $(2/3 + O(1/n))$ -MMS allocation always exists [36] and can be computed in polynomial time [2]. We use both EF1 and an approximation of the MMS guarantee as measures of fairness.

Additionally, we also want our solution to be economically efficient.² To this end, we use the rather unrestrictive notion of *Pareto optimality*.

Definition 2.4 (PO: Pareto Optimality). An allocation $\mathbf{A} \in \Pi_n(\mathcal{M})$ is called Pareto optimal if no alternative allocation $\mathbf{A}' \in \Pi_n(\mathcal{M})$ can make some players strictly better off without making any player strictly worse off. Formally, we require that

$$\forall \mathbf{A}' \in \Pi_n(\mathcal{M}), \left(\exists i \in \mathcal{N}, v_i(A'_i) > v_i(A_i) \right) \Rightarrow \left(\exists j \in \mathcal{N}, v_j(A'_j) < v_j(A_j) \right).$$

¹To be perfectly accurate, this is not satisfied if A_j is empty, but, clearly, in this case i does not envy j .

²In the absence of this requirement, even envy freeness can be achieved by simply not allocating any goods.

3 MAXIMUM NASH WELFARE IS EF1 AND PO

The gold standard of fairness—envy freeness (EF)—cannot be guaranteed in the context of indivisible goods. In contrast, envy freeness up to one good (EF1) is surprisingly easy to achieve under additive valuations.

Indeed, under the *draft* mechanism, the goods are allocated in a round-robin fashion: each of the players $1, \dots, n$ selects her most preferred good in that order, and we repeat this process until all the goods have been selected. To see why this allocation is EF1, consider some player $i \in \mathcal{N}$. We can partition the sequence of choices $1, \dots, i-1, i, i+1, \dots, n, 1, \dots, i-1, \dots$ into *phases* $i, \dots, i-1$, each starting when player i makes a choice, and ending just before she makes the next choice. In each phase, i receives a good that she (weakly) prefers to each of the $n-1$ goods selected by subsequent players. The only potential source of envy is the goods selected by players $1, \dots, i-1$ before the beginning of the first phase (that is, before i ever chose a good); but there is at most one such good per player $j \in [i-1]$, and removing that good from the bundle of j eliminates any envy that i might have had towards j .

However, it is clear that the allocation returned by the draft mechanism is not guaranteed to be Pareto optimal. One intuitive way to see this is that the draft outcome is highly constrained, in that all players receive almost the same number of goods; and mutually beneficial swaps of one good in return for multiple goods are possible.

Is there a different approach for generating allocations that are EF1 and PO? Surprisingly, several natural candidates fail. For example, maximizing the utilitarian welfare (the sum of utilities to the players) or the egalitarian welfare (the minimum utility to any player) is not EF1 (see Example B.2 in Appendix B). Interestingly, maximizing these objectives subject to the constraint that the allocation is EF1 violates PO (see Example B.3 in Appendix B, which was generated through computer simulations).

An especially promising idea—which was our starting point for the research reported herein—is to compute a CEEI allocation assuming the goods are divisible, and then to come up with an intelligent rounding scheme to allocate each good to one of the players who received some fraction of it. The hope was that, because the CEEI allocation is known to be EF for divisible goods [41], some rounding scheme, while inevitably violating EF, will only create envy up to one good, i.e., will still satisfy EF1. But, we found a counterexample in which *every* rounding of the “divisible CEEI” allocation violates EF1; this is presented as Example B.1 in Appendix B.

As mentioned earlier, for divisible goods, a CEEI allocation maximizes the Nash welfare. And, although a CEEI allocation may not exist for indivisible goods, one can still maximize the Nash welfare over all feasible allocations. Strikingly, this solution, which we refer to as the *maximum Nash welfare* (MNW) solution, achieves both EF1 and PO.

Definition 3.1 (The MNW solution). The Nash welfare of allocation $\mathbf{A} \in \Pi_n(\mathcal{M})$ is defined as $\text{NW}(\mathbf{A}) = \prod_{i \in \mathcal{N}} v_i(A_i)$. Given valuations $\{v_i\}_{i \in \mathcal{N}}$, the MNW solution selects an allocation $\mathbf{A}^{\text{MNW}}$ maximizing the Nash welfare among all feasible allocations, i.e.,

$$\mathbf{A}^{\text{MNW}} \in \arg \max_{\mathbf{A} \in \Pi_n(\mathcal{M})} \text{NW}(\mathbf{A}).$$

If it is possible to achieve positive Nash welfare (i.e., provide a positive utility to every player simultaneously), any Nash-welfare-maximizing allocation can be selected. The edge case in which every feasible allocation has zero Nash welfare (i.e., it is impossible to provide positive utility to every player simultaneously) is highly unlikely to appear in practice, but it must be handled carefully to retain the solution’s attractive fairness and efficiency properties.

In more detail, computing the MNW solution consists of two stages: (i) finding a largest set of players S to which one can simultaneously provide a positive utility (if there are multiple such sets

ALGORITHM 1: The MNW solution**Input:** The set of players $\mathcal{N}$, the set of indivisible goods $\mathcal{M}$, and players' valuations $\{v_i\}_{i \in \mathcal{N}}$ **Output:** An MNW allocation $\mathbf{A}^{\text{MNW}}$

- 1 $S \in \arg \max_{T \subseteq \mathcal{N}: \exists \mathbf{A} \in \Pi_n(\mathcal{M}), \forall i \in T, v_i(A_i) > 0} |T|$; // a largest set of players that can be
// simultaneously given a positive utility
- 2 $\mathbf{A}^* \leftarrow \arg \max_{\mathbf{A} \in \Pi_{|S|}(\mathcal{M})} \prod_{i \in S} v_i(A_i)$; // The MNW allocation to players in S
- 3 $A_i^{\text{MNW}} \leftarrow A_i^*, \forall i \in S$;
- 4 $A_i^{\text{MNW}} \leftarrow \emptyset, \forall i \in \mathcal{N} \setminus S$; // Players in $\mathcal{N} \setminus S$ do not receive any goods

S , our results hold independently of the tie-breaking), and (ii) finding an allocation of the goods to players in S that maximizes their product of utilities. The MNW solution is formally described in Algorithm 1. We say that an allocation is a *maximum Nash welfare (MNW)* allocation if it can be selected by the MNW solution.

We are now ready to state our first result, which is relatively simple yet, we believe, especially compelling.

THEOREM 3.2. *Every MNW allocation is envy-free up to one good (EF1) and Pareto optimal (PO) for additive valuations over indivisible goods.*

PROOF. Let $\mathbf{A}$ denote an MNW allocation. First, let us assume $\text{NW}(\mathbf{A}) > 0$. Pareto optimality of $\mathbf{A}$ holds trivially, because an alternative allocation that increases the utility to some players without decreasing the utility to any player would increase the Nash welfare, contradicting the optimality of the Nash welfare under $\mathbf{A}$. Suppose, for contradiction, that $\mathbf{A}$ is not EF1, and that player i envies player j even after removing any single good from player j 's bundle.

Let $g^* = \arg \min_{g \in A_j, v_i(g) > 0} v_j(g)/v_i(g)$. Note that g^* is well defined, because player i envying player j implies that player i has a positive value for at least one good in A_j . Let $\mathbf{A}'$ denote the allocation obtained by moving g^* from player j to player i in $\mathbf{A}$. We now show that $\text{NW}(\mathbf{A}') > \text{NW}(\mathbf{A})$, which gives the desired contradiction, as the Nash welfare is optimal under $\mathbf{A}$. Specifically, we show that $\text{NW}(\mathbf{A}')/\text{NW}(\mathbf{A}) > 1$. The ratio is well defined, because we assumed $\text{NW}(\mathbf{A}) > 0$.

Note that $v_k(A'_k) = v_k(A_k)$ for all $k \in \mathcal{N} \setminus \{i, j\}$, $v_i(A'_i) = v_i(A_i) + v_i(g^*)$, and $v_j(A'_j) = v_j(A_j) - v_j(g^*)$. Hence,

$$\frac{\text{NW}(\mathbf{A}')}{\text{NW}(\mathbf{A})} > 1 \Leftrightarrow \left[1 - \frac{v_j(g^*)}{v_j(A_j)}\right] \cdot \left[1 + \frac{v_i(g^*)}{v_i(A_i)}\right] > 1 \Leftrightarrow \frac{v_j(g^*)}{v_i(g^*)} \cdot [v_i(A_i) + v_i(g^*)] < v_j(A_j), \quad (1)$$

where the last transition follows using simple algebra. Due to our choice of g^* , we have

$$\frac{v_j(g^*)}{v_i(g^*)} \leq \frac{\sum_{g \in A_j} v_j(g)}{\sum_{g \in A_j} v_i(g)} = \frac{v_j(A_j)}{v_i(A_j)}. \quad (2)$$

Because player i envies player j even after removing g^* from player j 's bundle, we have

$$v_i(A_i) + v_i(g^*) < v_i(A_j). \quad (3)$$

Multiplying Equations (2) and (3) gives us the desired Equation (1).

Let us now address the special case where $\text{NW}(\mathbf{A}) = 0$. Let S denote the set of players to which the solution gives positive utility. Then, by the definition of the MNW solution (see Algorithm 1), S is a largest set of players to which one can provide positive utility. Pareto optimality of $\mathbf{A}$ now follows easily. An alternative allocation that does not reduce the utility to any player (and thus gives positive utility to each player in S) cannot give positive utility to any player in $\mathcal{N} \setminus S$. It also

cannot increase the utility to a player in S , because that would increase the product of utilities to the players in S , which $\mathbf{A}$ already maximizes.

From the proof of the case of $\text{NW}(\mathbf{A}) > 0$, we already know that there is no envy up to one good among players in S , because $\mathbf{A}$ is an MNW allocation over these players, and under $\mathbf{A}$ the product of utilities to the players in S is positive. Further, because players in $\mathcal{N} \setminus S$ do not receive any goods, we only need to show that player $i \in \mathcal{N} \setminus S$ does not envy player $j \in S$ up to one good. Suppose for contradiction that she does. Choose $g_j \in A_j$ such that $v_j(g_j) > 0$. Such a good exists, because we know $v_j(A_j) > 0$. Because player i envies player j up to one good, we have $v_i(A_j \setminus \{g_j\}) > v_i(A_i) = 0$. Hence, there exists a good $g_i \in A_j \setminus \{g_j\}$ such that $v_i(g_i) > 0$. However, in that case moving good g_i from player j to player i provides positive utility to player i while retaining positive utility to player j (because player j still has good g_j with $v_j(g_j) > 0$). This contradicts the fact that S is a largest set of players to which one can provide positive utility. Hence, the MNW allocation $\mathbf{A}$ is both EF1 and PO. $\square$

3.1 General Valuations

Heretofore, we have focused on the case of additive valuations. As we argued earlier, this case is crucial in practice. But it is nevertheless of theoretical interest to understand whether the guarantees extend to larger classes of combinatorial valuations.

Specifically, Theorem 3.2 states that MNW guarantees EF1 and PO under additive valuations. We ask whether EF1 and PO can be achieved simultaneously by any algorithm, not necessarily MNW, under subadditive, superadditive, submodular (a special case of subadditive), and supermodular (a special case of superadditive) valuations. The definitions of these valuation classes as well as the proofs of all the results in this section are provided in Appendix C. Unfortunately, we obtain a negative result for three of the four valuation classes.

THEOREM 3.3. *For the classes of subadditive and supermodular (and thus superadditive) valuations over indivisible goods, there exist instances that do not admit allocations that are envy-free up to one good and Pareto optimal.*

We were unable to settle this question for the class of submodular valuations. And although there exist examples with submodular valuations (see, e.g., Example C.3) in which no MNW allocation satisfies EF1, we can show that every MNW allocation satisfies a relaxation of EF1 together with PO.

Definition 3.4 (MEF1: Marginal Envy Freeness Up to One Good). We say that an allocation $\mathbf{A} \in \Pi_n(\mathcal{M})$ satisfies MEF1 if

$$\forall i, j \in \mathcal{N}, \exists g \in A_j, v_i(A_i) \geq v_i(A_i \cup A_j \setminus \{g\}) - v_i(A_i).$$

In comparing the definition of MEF1 to the definition of EF1, we see that on the right-hand side, $v_i(A_j \setminus \{g\})$, i.e., the value of player i for $A_j \setminus \{g\}$, is replaced by $v_i(A_i \cup A_j \setminus \{g\}) - v_i(A_i)$, which is the *marginal value* of player i for $A_j \setminus \{g\}$ given that player i is already allocated A_i . For submodular valuations, MEF1 is strictly weaker than EF1, while for additive valuations, MEF1 coincides with EF1. Hence, Theorem 3.2 follows directly from the next result (although our direct proof of Theorem 3.2 is simpler).

THEOREM 3.5. *Every MNW allocation satisfies marginal envy freeness up to one good (MEF1) and Pareto optimality (PO) for submodular valuations over indivisible goods.*

4 MAXIMUM NASH WELFARE IS APPROXIMATELY MMS

In this section, we show that the fairness properties of the MNW solution extend to an alternative relaxation of envy freeness—the maximin share guarantee, as well as a variant thereof—in theory and practice.

4.1 Approximate MMS, in Theory

From a technical viewpoint, our most involved result is the following theorem.

THEOREM 4.1. *Every MNW allocation is π_n -maximin share (MMS) for additive valuations over indivisible goods, where*

$$\pi_n = \frac{2}{1 + \sqrt{4n - 3}}.$$

Further, the factor π_n is tight, i.e., for every $n \in \mathbb{N}$ and $\epsilon > 0$, there exists an instance with n players having additive valuations in which no MNW allocation is $(\pi_n + \epsilon)$ -MMS.

Before we provide a proof, let us recall that the best-known approximation of the MMS guarantee—to date—is $2/3 + O(1/n)$ [36], where the bound for $n = 3$ is $3/4$. But the only known way to achieve a good bound is to build the algorithm around the MMS approximation goal [2, 36]. In contrast, the MNW solution achieves its $\pi_n = \Theta(1/\sqrt{n})$ ratio “organically” as one of several attractive properties. Moreover, in almost all real-world instances, the number of players n is fairly small. For example, on Spliddit, the average number of players is very close to 3, for which our worst-case approximation guarantee is $\pi_3 = 1/2$ —qualitatively similar to $3/4$. That said, the approximation ratio achieved on real-world instances is significantly better (see Section 4.3).

PROOF OF THEOREM 4.1. We first prove that an MNW allocation is π_n -MMS (lower bound), and later prove tightness of the approximation ratio π_n (upper bound).

Proof of the lower bound: Let $\mathbf{A}$ be an MNW allocation. As in the proof of Theorem 3.2, we begin by assuming $\text{NW}(\mathbf{A}) > 0$, and handle the case of $\text{NW}(\mathbf{A}) = 0$ later. Fix a player $i \in \mathcal{N}$. For a player $j \in \mathcal{N} \setminus \{i\}$, let $g_j^* = \arg \max_{g \in A_j} v_i(g)$ denote the good in player j ’s bundle that player i values the most. We need to establish an important lemma.

LEMMA 4.2. *It holds that*

$$v_i(A_j \setminus \{g_j^*\}) \leq \min \left\{ v_i(A_i), \frac{(v_i(A_i))^2}{v_i(g_j^*)} \right\},$$

where the RHS is defined to be $v_i(A_i)$ if $v_i(g_j^*) = 0$.

PROOF. First, $v_i(A_j \setminus \{g_j^*\}) \leq v_i(A_i)$ follows directly from the fact that $\mathbf{A}$ is an MNW allocation, and is therefore EF1 (Theorem 3.2). If $v_i(g_j^*) = 0$, then we are done. Assume $v_i(g_j^*) > 0$. By the definition of an MNW allocation, moving good g_j^* from player j to player i should not increase the Nash welfare. Thus,

$$v_i(A_i \cup \{g_j^*\}) \cdot v_j(A_j \setminus \{g_j^*\}) \leq v_i(A_i) \cdot v_j(A_j) \Rightarrow v_j(g_j^*) \geq v_j(A_j) - \frac{v_i(A_i) \cdot v_j(A_j)}{v_i(A_i \cup \{g_j^*\})}. \quad (4)$$

Note that the RHS in the above expression is positive, because $v_i(g_j^*) > 0$. Hence, we also have $v_j(g_j^*) > 0$. Similarly, moving all the goods in A_j except g_j^* from player j to player i should also not increase the Nash welfare. Hence,

$$v_i(A_i \cup A_j \setminus \{g_j^*\}) \cdot v_j(g_j^*) \leq v_i(A_i) \cdot v_j(A_j).$$

We conclude that

$$\begin{aligned} v_i(A_j \setminus \{g_j^*\}) &\leq \frac{v_i(A_i) \cdot v_j(A_j)}{v_j(g_j^*)} - v_i(A_i) \leq \frac{v_i(A_i) \cdot v_j(A_j)}{v_j(A_j) - \frac{v_i(A_i) \cdot v_j(A_j)}{v_i(A_i \cup \{g_j^*\})}} - v_i(A_i) \\ &= v_i(A_i) \cdot \left[\frac{1}{1 - \frac{v_i(A_i)}{v_i(A_i \cup \{g_j^*\})}} - 1 \right] = v_i(A_i) \cdot \left[\frac{v_i(A_i \cup \{g_j^*\})}{v_i(g_j^*)} - 1 \right] = \frac{(v_i(A_i))^2}{v_i(g_j^*)}, \end{aligned}$$

where the second transition follows from Equation (4). ■ (Proof of Lemma 4.2)

Now, let us find an upper bound on the MMS guarantee for player i . Recall that MMS_i is the maximum value player i can guarantee herself if she partitions the set of goods into n bundles but receives her least-valued bundle. The key intuition is that indivisibility of the goods only restricts the player in terms of the partitions she can create. That is, if some of the goods become *divisible*, it can only increase the MMS guarantee of the player, as she can still create all the bundles that she could with indivisible goods.

Suppose all the goods except goods in $T = \{g_j^* : j \in \mathcal{N} \setminus \{i\}, v_i(g_j^*) > \text{MMS}_i\}$ become divisible. It is easy to see that in the following partition, player i 's value for each bundle must be at least MMS_i : put each good in T (entirely) in its own bundle, and divide the rest of the goods into $n - |T|$ bundles of equal value to player i . Because each of the latter $n - |T|$ bundles must have value at least MMS_i for player i , we get

$$\text{MMS}_i \leq \frac{v_i(A_i) + \sum_{j \in \mathcal{N} \setminus \{i\}} (v_i(g_j^*) \cdot \mathbb{I}[v_i(g_j^*) \leq \text{MMS}_i] + v_i(A_j \setminus \{g_j^*\}))}{n - \sum_{j \in \mathcal{N} \setminus \{i\}} \mathbb{I}[v_i(g_j^*) > \text{MMS}_i]}, \quad (5)$$

where $\mathbb{I}(\cdot)$ denotes the indicator function.

Next, we use the upper bound on $v_i(A_j \setminus \{g_j^*\})$ from Lemma 4.2, and divide both sides of Equation (5) by $v_i(A_i)$. For simplicity, let us denote $x_j = v_i(g_j^*)/v_i(A_i)$, and $\beta = \text{MMS}_i/v_i(A_i)$. Note that β is the reciprocal of the bound on the MMS approximation that we are interested in. Then, we get

$$\beta \leq \frac{1 + \sum_{j \in \mathcal{N} \setminus \{i\}} \left(x_j \cdot \mathbb{I}[x_j \leq \beta] + \min \left\{ 1, \frac{1}{x_j} \right\} \right)}{n - \sum_{j \in \mathcal{N} \setminus \{i\}} \mathbb{I}[x_j > \beta]}.$$

Let $f(\mathbf{x}; \beta)$ denote the RHS of the inequality above. Then, we can write $\beta \leq f(\mathbf{x}; \beta) \leq \max_{\mathbf{x}} f(\mathbf{x}; \beta)$. Note that if $\beta \leq 1$, then player i is already receiving her full maximin share value, which gives a (stronger than) desired MMS approximation. Let us therefore assume that $\beta > 1$. To find the maximum value of $f(\mathbf{x}; \beta)$ over all $\mathbf{x}$, let us take its partial derivative with respect to x_k for $k \in \mathcal{N} \setminus \{i\}$. Note that the function is differentiable at all points except $x_k = 1$ and $x_k = \beta$.

$$\frac{\partial f}{\partial x_k} = \begin{cases} \frac{1}{n - \sum_{j \in \mathcal{N} \setminus \{i\}} \mathbb{I}[x_j > \beta]} & \text{if } 0 \leq x_k < 1, \\ \frac{1 - (x_k)^{-2}}{n - \sum_{j \in \mathcal{N} \setminus \{i\}} \mathbb{I}[x_j > \beta]} & \text{if } 1 < x_k < \beta, \\ \frac{-(x_k)^{-2}}{n - \sum_{j \in \mathcal{N} \setminus \{i\}} \mathbb{I}[x_j > \beta]} & \text{if } \beta < x_k. \end{cases}$$

Note that $\partial f / \partial x_k > 0$ for $x \in (0, 1)$ and $x \in (1, \beta)$, and $\partial f / \partial x_k < 0$ for $x_k > \beta$. Further note that f is continuous at $x_k = 1$. Hence, the maximum value of f is achieved either at $x_k = \beta$ or in the limit as $x_k \rightarrow \beta^+$ (i.e., when x_k converges to β from above). Suppose the maximum is achieved when t

of the x_k 's are equal to β , and the other $n - t - 1$ approach β from above. Then, the value of f is

$$g(t; \beta) = \frac{1 + t \cdot \left(\beta + \frac{1}{\beta}\right) + (n - t - 1) \cdot \frac{1}{\beta}}{n - (n - t - 1)}.$$

We now have that $\beta \leq \max_{t \in \{0, \dots, n-1\}} g(t; \beta)$. Note that

$$\frac{\partial g}{\partial t} = \frac{\beta - 1 - (n - 1) \cdot \frac{1}{\beta}}{(t + 1)^2}.$$

If $\beta = \text{MMS}_i / v_i(A_i) \leq 1/\pi_n$, we already have the desired MMS approximation. Assume $\beta > 1/\pi_n$. It is easy to check that this implies $\partial g / \partial t > 0$. Thus, the maximum value of g is achieved at $t = n - 1$, which gives $\beta \leq (1/n) \cdot (1 + (n - 1) \cdot (\beta + 1/\beta))$, which simplifies to $\beta \leq 1/\pi_n$, which is a contradiction as we assumed $\beta > 1/\pi_n$.

Recall that for the proof above, we assumed $\text{NW}(\mathbf{A}) > 0$. Let us now handle the special case where an MNW allocation $\mathbf{A}$ satisfies $\text{NW}(\mathbf{A}) = 0$. Let S denote the set of players that receive positive utility under $\mathbf{A}$, where $|S| < n$. Due to the definition of an MNW allocation (see Algorithm 1), $\mathbf{A}$ is an MNW allocation over the players in S . Thus, from the proof of the previous case, we know that each player in S in fact achieves at least a $\pi_{|S|}$ -fraction of her $|S|$ -player MMS guarantee, which is at least a π_n -fraction of her n -player MMS guarantee. Players in $\mathcal{N} \setminus S$ receive zero utility. We now show that their (n -player) MMS guarantee is also 0, which yields the required result.

Suppose a player $i \in \mathcal{N} \setminus S$ has a positive value for at least n goods in $\mathcal{M}$. Now, because these goods are allocated to at most $n - 1$ players in S , at least one player $j \in S$ must have received at least two goods g_1 and g_2 , both of which player i values positively. Because player j receives positive utility under $\mathbf{A}$ (i.e., $v_j(A_j) > 0$), it is easy to check that there exists a good $g \in \{g_1, g_2\}$ such that $v_j(A_j \setminus \{g\}) > 0$. Thus, moving good g to player i provides positive utility to player i while retaining positive utility to player j , which violates the fact that S is a largest set of players to that one can simultaneously provide positive utility. This shows that player i has positive utility for at most $n - 1$ goods in $\mathcal{M}$, which immediately implies $\text{MMS}_i = 0$, as required.

Proof of the upper bound (tightness): We now show that for every $n \in \mathbb{N}$ and $\epsilon > 0$, there exists an instance with n players in which no MNW allocation is $(\pi_n + \epsilon)$ -MMS. For $n = 1$, this is trivial, because $\pi_1 = 1$. Hence, assume $n \geq 2$.

Let the set of players be $\mathcal{N} = \{1, \dots, n\}$, and the set of goods be $\mathcal{M} = \{x\} \cup \bigcup_{i \in \{2, \dots, n\}} \{h_i, l_i\}$. Thus, we have $m = 2n - 1$ goods. We refer to h_i 's as the “heavy” goods and l_i 's as the “light” goods. Let the valuations of the players for the goods be as follows: Choose a sufficiently small $\epsilon' > 0$ (an upper bound on ϵ' will be determined later in the proof).

$$\text{Player 1: } v_1(x) = 1, \text{ and } \forall j \in \{2, \dots, n\}, v_1(h_j) = \frac{1}{\pi_n} - \epsilon' \text{ and } v_1(l_j) = \pi_n - \epsilon'.$$

$$\text{Player } i, \text{ for } i \geq 2: v_i(h_i) = \frac{1}{\pi_n + 1}, v_i(l_i) = \frac{\pi_n}{\pi_n + 1}, \text{ and } \forall g \in \mathcal{M} \setminus \{h_i, l_i\}, v_i(g) = 0.$$

In particular, note that player 1 has a positive value for every good (for $\epsilon' < \pi_n$), while for $i \geq 2$, player i has a positive value for only two goods: h_i and l_i . Consider the allocation $\mathbf{A}^*$ that assigns good x to player 1, and for every $i \in \mathcal{N} \setminus \{1\}$, allocates goods h_i and l_i to player i . We claim that $\mathbf{A}^*$ is the unique MNW allocation but is not $(\pi_n + \epsilon)$ -MMS.

First, note that an MNW allocation is Pareto optimal, and therefore it must allocate good x to player 1, because no other player has a positive value for x . Further, $\text{NW}(\mathbf{A}^*) > 0$, which implies that every MNW allocation must also have a positive Nash welfare. This in turn implies that an MNW allocation must assign to each player in $\mathcal{N} \setminus \{1\}$ at least one of h_i and l_i . Subject to these constraints, consider a candidate allocation $\mathbf{A}$.

Let p (resp. q) denote the number of players $i \in \mathcal{N} \setminus \{1\}$ that only receive good h_i (resp. l_i), and have utility $1/(\pi_n + 1)$ (resp. $\pi_n/(\pi_n + 1)$). Hence, exactly $n - 1 - p - q$ players $i \in \mathcal{N} \setminus \{1\}$ receive both h_i and l_i , and have utility 1. Player 1 receives good x , q heavy goods, and p light goods, and has utility $1 + q \cdot (1/\pi_n - \epsilon') + p \cdot (\pi_n - \epsilon')$. Thus, the Nash welfare of $\mathbf{A}$ is given by

$$\left(1 + q \cdot \left(\frac{1}{\pi_n} - \epsilon'\right) + p \cdot (\pi_n - \epsilon')\right) \left(\frac{1}{\pi_n + 1}\right)^p \left(\frac{\pi_n}{\pi_n + 1}\right)^q = \frac{1 + q \cdot \left(\frac{1}{\pi_n} - \epsilon'\right) + p \cdot (\pi_n - \epsilon')}{(1 + \pi_n)^p \cdot \left(1 + \frac{1}{\pi_n}\right)^q}.$$

Using binomial expansion, it is easy to show that the denominator in the final expression above is at least $1 + p \cdot \pi_n + q/\pi_n$, which is never less than the numerator, and is equal to the numerator if and only if $p = q = 0$. Note that $p = q = 0$ indeed gives our desired allocation $\mathbf{A}^*$. Hence, the maximum Nash welfare of 1 is uniquely achieved by the allocation $\mathbf{A}^*$.

Next, let us analyze the MMS guarantee for player 1. In particular, consider the partition of the set of goods into n bundles $B_1, \dots, B_n$ such that $B_1 = \{x, l_2, \dots, l_n\}$ and $B_i = \{h_i\}$ for all $i \in \{2, \dots, n\}$. Note that for all $i \in \{2, \dots, n\}$, $v_1(B_i) = 1/\pi_n - \epsilon'$. Also,

$$v_1(B_1) = 1 + (n - 1) \cdot (\pi_n - \epsilon') = 1 + (n - 1) \cdot \pi_n - (n - 1) \cdot \epsilon' = \frac{1}{\pi_n} - (n - 1) \cdot \epsilon',$$

where the final equality holds, because π_n is chosen precisely to satisfy the equation $1 + (n - 1) \cdot \pi_n = 1/\pi_n$. As the MMS guarantee of player 1 is at least her minimum value for any bundle in $\{B_1, \dots, B_n\}$, we have $\text{MMS}_1 \geq 1/\pi_n - (n - 1) \cdot \epsilon'$. In contrast, under the MNW allocation $\mathbf{A}^*$, we have $v_1(A_1) = 1$. Thus, the MMS approximation ratio on this instance is at most $1/(1/\pi_n - (n - 1) \cdot \epsilon')$. It is easy to check that for driving this ratio below $\pi_n + \epsilon$, it is sufficient to set

$$\epsilon' < \min \left\{ \pi_n, \frac{\epsilon}{(n - 1) \cdot \pi_n \cdot (\pi_n + \epsilon)} \right\}.$$

This completes the entire proof. ■ (Proof of Theorem 4.1)

A striking aspect of the proof of Theorem 4.1 is that, at first glance, the lower bound of π_n seems very loose. For example, key steps in the proof involve the derivation of an upper bound on the MMS guarantee of player i by assuming that some of the goods are divisible, and the maximization of the function $f(\cdot)$ over an unrestricted domain. Yet the ratio π_n turns out to be completely tight.

4.2 Approximate Pairwise MMS, in Theory

Adding to the conceptual arguments in favor of Theorem 4.1 (see the discussion just after the theorem statement), we note that it also has some rather striking implications. Let us first define a novel fairness property:

Definition 4.3 (α -Pairwise Maximin Share Guarantee). We say that an allocation $\mathbf{A} \in \Pi_n(\mathcal{M})$ is an α -pairwise maximin share (MMS) allocation if

$$\forall i, j \in \mathcal{N}, v_i(A_i) \geq \alpha \cdot \max_{\mathbf{B} \in \Pi_2(A_i \cup A_j)} \min\{v_i(B_1), v_i(B_2)\}.$$

We simply say that $\mathbf{A}$ is pairwise MMS if it is 1-pairwise MMS. Note that the pairwise MMS guarantee is similar to the MMS guarantee, but instead of player i partitioning the set of all items into n bundles, she partitions the combined bundle of herself and another player into two bundles, and receives the one she values less. Although the example below shows that neither the pairwise MMS guarantee nor the MMS guarantee implies the other, we show in Theorem 4.6 that a pairwise MMS allocation is $(1/2)$ -MMS.

Example 4.4 (Neither pairwise MMS nor MMS implies the other). We show that neither pairwise MMS nor MMS implies the other even when players have identical valuations.

Showing that MMS does not imply pairwise MMS is easy. In fact, we can use our motivating example from the introduction. Consider a set of three players $N = \{1, 2, 3\}$, and a set of five goods M . Let each player have value 1 for each good. Consider the allocation A that assigns three goods to player 1, and a single good to each of players 2 and 3. Because the maximin share of each player is 1, A is an MMS allocation. However, it is easy to check that A violates pairwise MMS.

Next, we show that pairwise MMS does not imply MMS. Consider a set of three players $N = \{1, 2, 3\}$ and a set of seven goods $M = \{g_i : i \in [7]\}$. Let the players have an identical valuation function v , where $v(g_1) = 1$, $v(g_2) = v(g_3) = 2$, $v(g_4) = v(g_5) = 3$, $v(g_6) = 4$, and $v(g_7) = 6$.

Consider the allocation A where $A_1 = \{g_7\}$, $A_2 = \{g_3, g_4, g_5\}$, and $A_3 = \{g_1, g_2, g_6\}$. The values derived by the three players are 6, 8, and 7, respectively, while the MMS guarantee of each player is 7, because the goods can be partitioned into three sets of value 7 each. Thus, A is not an MMS allocation, because the MMS guarantee of player 1 is violated. It is easy to check that A is nonetheless a pairwise MMS allocation.

We do not know whether a pairwise MMS allocation always exists (under the constraint that all goods must be allocated). In fact, there is an even more tantalizing and elusive fairness notion that is strictly weaker than pairwise MMS, but strictly stronger than EF1 (see Theorem 4.6 below, which, in particular, implies that pairwise MMS is stronger than EF1).

Definition 4.5 (EFX: Envy freeness Up to the Least-valued Good). We say that an allocation $A \in \Pi_n(M)$ is envy-free up to the least (positively) valued good if

$$\forall i, j \in N, \forall g \in A_j \text{ such that } v_i(g) > 0, v_i(A_i) \geq v_i(A_j \setminus \{g\}).$$

While EF1 requires that player i not envy player j after the removal of player i 's *most valued* good from player j 's bundle, EFX requires that this no-envy condition would hold even after the removal of player i 's *least positively valued* good from player j 's bundle. Despite significant effort, we were not able to settle the question of whether an EFX allocation always exists (assuming all goods must be allocated), and leave it as an enigmatic open question.

At this point, the reader may be wondering about the abundance of fairness notions we are considering. But they are all related, as the following result shows.

THEOREM 4.6. *For additive valuations over indivisible goods, the pairwise maximin share guarantee is implied by envy freeness (EF), and implies $1/2$ -maximin share guarantee, envy freeness up to the least-valued good (EFX), and as a direct consequence, envy freeness up to one good (EF1).*

PROOF. Let A be an EF allocation, i.e., $v_i(A_i) \geq v_i(A_j)$ for all pairs of players $i, j \in N$. Let PMMS_i denote the pairwise MMS guarantee of player i :

$$\text{PMMS}_i = \max_{j \in N \setminus \{i\}} \max_{B \in \Pi_2(A_i \cup A_j)} \min\{v_i(B_1), v_i(B_2)\}.$$

Then, we have

$$\text{PMMS}_i \leq \max_{j \in N \setminus \{i\}} \frac{v_i(A_i) + v_i(A_j)}{2} \leq v_i(A_i),$$

where the first transition holds, because its right-hand side is the pairwise MMS guarantee that player i would have if all goods were divisible, which is an upper bound on PMMS_i , because divisible goods offer the player a greater flexibility in partitioning the goods. The second transition follows directly from the envy freeness of A .

Next, let A be a pairwise MMS allocation. It is easy to show that A must also be EFX: if player i envies player j after the removal of player i 's least positively valued good g^* from A_j , then it follows that player i 's pairwise MMS guarantee is at least $v_i(A_i \cup \{g^*\}) > v_i(A_i)$ due to the partition

$(A_i \cup \{g^*\}, A_j \setminus \{g^*\})$. However, this implies that $\mathbf{A}$ is not pairwise MMS, which is a contradiction. Hence, $\mathbf{A}$ is also EFX. It is trivial to check that EFX implies EF1 by definition; hence, $\mathbf{A}$ is also EF1.

Finally, we show that a pairwise MMS allocation $\mathbf{A}$ is also $1/2$ -MMS. Consider players i and j . There are only two possible cases: (i) A_j has at most one good that player i values positively, i.e., $|A_j \cap \{g \in \mathcal{M} \mid v_i(g) > 0\}| \leq 1$, or (ii) $v_i(A_j) \leq 2 \cdot v_i(A_i)$. Indeed, if A_j has at least two goods that player i values positively, and $v_i(A_j) > 2 \cdot v_i(A_i)$, then consider the good g^* that is the least valuable among player i 's positively valued goods in A_j . In that case, player i could partition $A_i \cup A_j$ into $(A_i \cup \{g^*\}, A_j \setminus \{g^*\})$ and ensure that her pairwise MMS value is strictly more than $v_i(A_i)$, which is a contradiction, because $\mathbf{A}$ is pairwise MMS.

Now, if no player in $\mathcal{N} \setminus \{i\}$ falls into case (ii), then it is easy to see that the MMS guarantee of player i is at most $v_i(A_i)$. If a non-empty subset $S \subseteq \mathcal{N} \setminus \{i\}$ of players fall into case (ii), then we can bound the MMS guarantee of player i from above by assuming that all goods allocated to players $S \cup \{i\}$ are divisible. However, this still gives an MMS guarantee of at most $2 \cdot v_i(A_i)$, because each player in $j \in S \cup \{i\}$ satisfies $v_i(A_j) \leq 2 \cdot v_i(A_i)$. Thus, the MMS guarantee of player i is at most $2 \cdot v_i(A_i)$, which implies that $\mathbf{A}$ is $1/2$ -MMS. $\square$

Given this backdrop for the pairwise MMS notion, it is interesting that our next result directly translates the MMS approximation bound of Theorem 4.1 into a pairwise MMS approximation.

COROLLARY 4.7. *Every MNW allocation is Φ -pairwise MMS, where Φ is the golden ratio conjugate, i.e., $\Phi = (\sqrt{5} - 1)/2 \approx 0.618$. Further, the factor Φ is tight, i.e., for every $n \in \mathbb{N}$ and $\epsilon > 0$, there exists an instance with n players having additive valuations in which no MNW allocation is $(\Phi + \epsilon)$ -pairwise MMS.*

PROOF. An MNW allocation $\mathbf{A}$ has the following interesting property: Take the goods allocated to players i and j , i.e., $\mathcal{M}' = A_i \cup A_j$, and take the set of players $\mathcal{N}' = \{i, j\}$. Then the allocation given by A_i and A_j is also an MNW allocation for the reduced instance of allocating the set of goods $\mathcal{M}'$ to the set of players $\mathcal{N}'$. This fact is easy to see when either $v_i(A_i) > 0$ and $v_j(A_j) > 0$ (otherwise, we could achieve higher Nash welfare), or $v_i(A_i) = v_j(A_j) = 0$. When $v_i(A_i) = 0$ but $v_j(A_j) > 0$ (without loss of generality), every allocation of $\mathcal{M}'$ to players $\{i, j\}$ must provide zero utility to at least one player, otherwise this part of the allocation could be used in the original instance to increase the number of players that receive positive utility, contradicting the fact that an MNW allocation provides positive utility to the maximum number of players. Hence, the allocation in the reduced instance that provides all the goods in $\mathcal{M}'$ to player j (which is exactly allocation $\mathbf{A}$ restricted to the reduced instance) is indeed an MNW allocation, and is π_2 -MMS in the reduced instance (Theorem 4.1).

We therefore conclude that the MNW allocation $\mathbf{A}$ is Φ -pairwise MMS in the original instance as $\pi_2 = \Phi$. To establish tightness of the factor Φ , for a given $n \in \mathbb{N}$ and $\epsilon > 0$, we simply use the example from the proof of the upper bound in Theorem 4.1 after replacing π_n by $\pi_2 = \Phi$ in the valuations of the players. In the new example, now the pairwise MMS approximation ratio (instead of the MMS approximation ratio in the original example) can be driven below $\pi_2 + \epsilon$ for a value of ϵ' less than $\min(\pi_2, \epsilon/(\pi_2 \cdot (\pi_2 + \epsilon)))$, which is a bound obtained by substituting $n = 2$ in the upper bound on ϵ' from the proof of Theorem 4.1. $\square$

4.3 Approximate MMS and Pairwise MMS, in Practice

Theorem 4.1 and Corollary 4.7 show that the MNW solution is *guaranteed* to be π_n -MMS and Φ -pairwise MMS. We now evaluate these approximation ratios in practice using real-world data. Specifically, we use the 1,281 instances created so far through Spliddit's "divide goods" application. The number of players in these instances ranges from 2 to 10, and the number of goods ranges from

Fig. 1. MMS and Pairwise MMS approximation of the MNW solution on real-world data from Spliddit.

3 to 93. Figures 1(a) and 1(b) show the histograms of the MMS and pairwise MMS approximation ratios, respectively, achieved by the MNW solution on these instances.

Most importantly, observe that the MNW solution provides every player her full MMS (resp. pairwise MMS) guarantee, i.e., achieves the ideal 1-approximation, in more than 95% (resp. 90%) of the instances. Further, in contrast to the tight worst-case ratios of $\pi_n = \Theta(1/\sqrt{n})$ and $\Phi \approx 0.618$, the MNW solution achieves a ratio of at least $3/4$ for both properties on *all* the real-world instances in our dataset.

5 IMPLEMENTATION

It is known that computing an exact MNW allocation is $\mathcal{NP}$ -hard even for 2 players with identical additive valuations, due to a simple reduction from the $\mathcal{NP}$ -hard problem PARTITION [33, 37]. Our goal in this section is to develop a fast implementation of the MNW solution, despite this obstacle. An existing approach to maximizing the Nash welfare [34] iteratively modifies an initial allocation to improve the Nash welfare at each step, but may return a local maximum that does not provide any fairness or efficiency guarantees. Instead, we use integer programming to find the global optimum in a scalable way. Note that most real-world instances are relatively small, but response time can be crucial. For example, Spliddit has a demo mode, where users expect almost instantaneous results. Moreover, some instances are actually very large, as we discuss below.

Let us begin by recalling that the first step in computing an MNW allocation is to find a largest set of players S that can be given positive utility simultaneously. For submodular valuations (and hence, for additive valuations) it holds that if a player has a positive value for a bundle of goods, there must exist a good g in the bundle such that the player has a positive value for the singleton set $\{g\}$. Thus, at least for submodular valuations, to provide a positive utility to the maximum number of players it is sufficient to restrict our attention to allocations that assign at most one good to each player. We create a bipartite graph G with the players on one side and the goods on the other, and add an edge from player i to good g iff $v_i(g) > 0$.³ Our desired set S can now be computed as the set of players satisfied under a maximum cardinality matching in G . There are many popular polynomial time algorithms that one can use to find a maximum cardinality matching in a bipartite graph, e.g., the Hopcroft-Karp method. While this shows that set S can be computed in polynomial time for submodular (and thus for additive) valuations, the problem may be computationally hard for other classes of valuation functions.

³Recall that $v_i(g)$ is shorthand for $v_i(\{g\})$.

$$\begin{aligned}
& \text{Maximize } \sum_{i \in \mathcal{N}} \log \left(\sum_{g \in \mathcal{M}} x_{i,g} \cdot v_i(g) \right) \\
& \text{subject to } \sum_{i \in \mathcal{N}} x_{i,g} = 1, \forall g \in \mathcal{M} \\
& \quad x_{i,g} \in \{0, 1\}, \forall i \in \mathcal{N}, g \in \mathcal{M}.
\end{aligned}$$

Fig. 2. Nonlinear discrete optimization.

$$\begin{aligned}
& \text{Maximize } \sum_{i \in \mathcal{N}} W_i \\
& \text{subject to } W_i \leq \log k + \lceil \log(k+1) - \log k \rceil \\
& \quad \times \left\lceil \sum_{g \in \mathcal{M}} x_{i,g} \cdot v_i(g) - k \right\rceil \\
& \quad \forall i \in \mathcal{N}, k \in \{1, 3, \dots, 999\} \\
& \quad \sum_{g \in \mathcal{M}} x_{i,g} \cdot v_i(g) \geq 1, \quad \forall i \in \mathcal{N} \\
& \quad \sum_{i \in \mathcal{N}} x_{i,g} = 1, \quad \forall g \in \mathcal{M} \\
& \quad x_{i,g} \in \{0, 1\}, \quad \forall i \in \mathcal{N}, g \in \mathcal{M}.
\end{aligned}$$

Fig. 4. MILP using segments.

Fig. 3. Logarithm and its approximations.

Fig. 5. Running time.

Once we find the set S , the task at hand reduces to computing an MNW allocation to the players in S . Hereinafter, we focus on this reduced problem. Thus, without loss of generality, we can assume that for the given set of players $\mathcal{N}$, an MNW allocation will achieve positive Nash welfare.

Figure 2 shows a simple mathematical program for computing an MNW allocation. The binary variable $x_{i,g}$ denotes whether player i receives good g . Subject to feasibility constraints, the program maximizes the sum of log of players' utilities, or, equivalently, the Nash welfare. Note that this is a discrete optimization program with a nonlinear objective, which is typically very hard to solve.

Fortunately, we can leverage some additional properties of the problem that arise in practice. Specifically, on Spliddit, users are required to submit integral additive valuations by dividing 1K points among the goods. This in turn ensures that the utilities to the players will also be integral, and not more than 1K. In theory, this does not help us: due to a known reduction from a strongly $\mathcal{NP}$ -complete problem—Exact Cover by 3-Sets (X3C)—to the problem of computing an MNW allocation [33], we cannot hope for a pseudopolynomial-time algorithm (i.e., a polynomial-time algorithm for Spliddit-like valuations). In practice, however, this structure of the valuations can be leveraged to convert the non-linear objective into a linear objective: $\sum_{i \in \mathcal{N}} \sum_{t=2}^{1000} (\log t - \log(t-1)) \cdot U_{i,t}$, where $U_{i,t} = \mathbb{I}[\sum_{g \in \mathcal{M}} x_{i,g} \cdot v_i(g) \geq t]$ for player $i \in \mathcal{N}$ and $t \in [1000]$ is an additional variable that can be encoded using two linear constraints. However, the introduction of $1000 \cdot n$ additional binary variables makes this approach impractical even for fairly small instances.

We therefore propose an alternative approach that introduces merely n continuous variables and, crucially, no integral variables. The trick is to use a continuous variable W_i denoting the log of the utility to player i , and bound it from above using a set of linear constraints such that the tightest bound at every integral point k is exactly $\log k$. This essentially replaces the log by a piecewise linear approximation thereof that has zero error at integral points. Figure 3 shows two such approximations of the log function (the red line): one that uses the tangent to the log curve at the point $(k, \log k)$ for each $k \in [1000]$ (the blue lines), and one that uses segments connecting points $(k, \log k)$ and $(k+1, \log(k+1))$ for each $k \in \{1, 3, \dots, 999\}$ (the green line). Each tangent and each segment is guaranteed to be an upper bound on the log function at every integral point due to the

concavity of $\log$.⁴ Importantly, note that the tightest upper bound at each *positive* integral point k is $\log k$. These transformations do not work at $k = 0$, i.e., they do not ensure $W_i = -\infty$ if player i gets zero utility. However, recall that in our subproblem each player can achieve a positive utility. Hence, we eliminate this concern by adding the constraints that each player must receive value of at least 1. We employ the transformation that uses segments, as it requires half as many constraints (and, incidentally, runs nearly twice as fast). Figure 4 shows the final mixed-integer linear program (MILP) with only n continuous and $n \cdot m$ binary variables, which is key to the practicability of this approach.

To assess how scalable our implementation is, we measure its running time on uniformly random Spliddit-like valuations; that is, uniformly random integral valuations that sum to 1K. We vary the number of players n from 5 to 50 in increments of 5, and keep the number of goods at $m = 3 \cdot n$ to match data from Spliddit, in which $m/n \approx 3$ on average. The experiments were performed on a 2.9GHz quad-core computer with 32GB RAM, using CPLEX to solve the MILPs. The indicator-variables-based approach failed to run within our time limit (60s) even for 5 players. Figure 5 shows the running time (averaged over 100 simulations, with the 5th and 95th percentiles) of the MILP formulation from Figure 4. Satisfyingly, we can solve instances with 50 players in less than 30s (whereas even the largest of the 1,281 instances on Spliddit has 10 players). In fact, the algorithm solves every Spliddit instance in less than 3s.

The largest real-world instance we have seen was actually reported offline by a Spliddit user. He needed to split an inheritance of roughly 1,400 goods with his 9 siblings. Our implementation solves an instance of this size in roughly 15s.

5.1 Precision Requirements

As our optimization program involves real-valued quantities (e.g., the logarithms), we must carefully set the precision level such that the optimal allocation computed up to the precision is *guaranteed* to be an MNW allocation. This is because an allocation that only approximately maximizes the Nash welfare may fail to satisfy the theoretical guarantees of an MNW allocation (Theorems 3.2 and 4.1, and Corollary 4.7).

Recall that our objective function is the log of the Nash welfare. Hence, the difference between the objective values of an (optimal) MNW allocation and any suboptimal allocation is at least $\log(1000^n) - \log(1000^n - 1) \geq 1/1000^n - (1/2)/1000^{2n}$, which can be captured using $O(n)$ bits of precision. This simple observation can be easily formalized to show that there exists $p \in O(n)$ such that if all the coefficients in the optimization program are computed up to p bits, and if the program is solved with p bits of precision (i.e., with an absolute error of at most 2^{-p} in the objective function), then the solution returned will indeed correspond to an MNW allocation. Crucially, p is independent of the number of goods. We expect the number of players n to be fairly small in everyday fair-division problems. For example, as previously mentioned, on Spliddit more than 95% of the instances for allocating indivisible goods have $n \leq 3$.

Nonetheless, if one's goal is solely to find an allocation that is EF1 and PO, a constant number of bits of precision would suffice. This is because capturing differences in objective values that are at least $\log(1000^2) - \log(1000^2 - 1)$ —a constant—ensures that the resulting allocation is EF1 and PO, as we show below.

- (1) *EF1*: Suppose the allocation is not EF1, and player i envies player j even after the removal of any single good from player j 's bundle. Then, our proof of Theorem 3.2 shows that we can

⁴In fact, this transformation is useful in maximizing any concave function, or minimizing any convex function, and thus may be of independent interest.

increase the Nash welfare by moving a specific good from player j to player i . Because this operation does not alter the utilities to all but two players, it must increase the logarithm of the Nash welfare by at least $\log(1000^2) - \log(1000^2 - 1)$, which is a contradiction, because our sensitivity level is sufficient to find this improvement.

- (2) *PO*: Suppose the allocation is not *PO*. Then there exists an alternative allocation that increases the utility to at least one player without decreasing the utility to any player. This must increase the logarithm of the Nash welfare by at least $\log(1000) - \log(1000 - 1) \geq \log(1000^2) - \log(1000^2 - 1)$, which is again a contradiction, because our sensitivity level is sufficient to find this improvement.

6 DISCUSSION

The goal of this article is to advocate for the Maximum Nash Welfare (MNW) solution for the fair allocation of goods. While it is justified by elegant fairness (EF1) and efficiency (PO) properties, these properties are not “sufficient” *in and of themselves*—they may allow undesirable outcomes (see Example B.4 in Appendix B). What makes the MNW solution compelling is that it provides intuitively fair outcomes, yet *organically* satisfies these formal fairness properties. Moreover, the MNW solution provides a $\Theta(1/\sqrt{n})$ -approximation to the MMS guarantee (Theorem 4.1), whereas an arbitrary EF1 and PO allocation only provides a $1/n$ -approximation (Theorem B.5 in Appendix B).

Throughout the article, we assumed that the goods are indivisible, but our results directly extend to the case where we have a mix of divisible and indivisible goods. The MNW solution in this case can be seen as the limit of the MNW solution on the instance where each divisible good is partitioned into k indivisible goods, as k goes to infinity. Theorem 3.2 therefore implies that the MNW solution is envy-free up to one *indivisible* good; that is, player i would not envy player j (who may have both divisible and indivisible goods) if one indivisible good is removed from the bundle of j . This provides an alternative proof for envy freeness of the MNW/CEEI solution when all goods are divisible. The results of Section 4 also directly go through—in fact, the proof of the MMS approximation result (Theorem 4.1) already “liquidates” some of the goods as a technical tool. Appendix A outlines the modified and scalable version of the implementation described in Section 5, which we have deployed on Spliddit, that can allocate a mix of divisible and indivisible goods.

It is remarkable that when all goods are divisible, three seemingly distinct solution concepts—the MNW solution, the CEEI solution, and proportional fairness (PF)—coincide. This is certainly not the case for indivisible goods: while a CEEI solution and a PF solution may not exist, the MNW solution always does. Nonetheless, our investigation revealed that even for indivisible goods, the PF solution and the MNW solution are closely related via a spectrum of solutions, which offers two advantages. First, it allows us to view the MNW solution as the optimal solution among those that lie on this spectrum and are guaranteed to exist. Second, it also gives a way to break ties—possibly even choose a unique allocation—among all MNW allocations. See Appendix D for a detailed analysis. This connection between MNW and PF raises an interesting question: Is it possible to relate the MNW solution to the CEEI solution when the goods are indivisible?

Finally, we have not addressed game-theoretic questions regarding the manipulability of the MNW solution. The reason is twofold. First, there are strong impossibility results that rule out reasonable strategy-proof solutions. For example, Schummer [38] shows that the only strategy-proof and Pareto optimal solutions are dictatorial—which means they are maximally unfair, if you will—even when there are only two players with linear utilities over divisible goods; clearly a similar result holds for indivisible goods (at least in an approximate sense).⁵ Second, we do not

⁵In theory, one can hope to circumvent this result by making manipulation computationally hard [7]. This is almost surely true (in the worst-case sense of hardness) for the MNW solution, where even computing the outcome is hard.

view manipulation as a major issue on Spliddit, because users are not fully aware of each other's preferences (they submit their evaluations in private), and—presumably, in most cases—have a very partial understanding of how the algorithm works.

APPENDIX

A IMPLEMENTATION ON SPLIDDIT

Section 5 outlines an implementation of the MNW solution when all the goods are indivisible. In contrast, our fair-division website Spliddit allows an arbitrary mix of divisible and indivisible goods, for which we designed an implementation that builds on the implementation of Section 5.

Splitting divisible goods:

As described in Section 6, one approach is to split each divisible good into k identical indivisible goods, and apply the MNW solution on the resulting set of indivisible goods. When k goes to infinity, this approach perfectly simulates the divisible goods, and gives the following relaxation of EF in addition to Pareto optimality (PO):

For every pair of players i and j , there exists an indivisible good in player j 's bundle such that player i does not envy player j after removing it from player j 's bundle.

However, splitting each divisible good into infinitely many indivisible goods is computationally not feasible. In practice, it suffices to split each divisible good into 100 indivisible goods, which provides the following relaxation of EF in addition to PO:

For every pair of players i and j , there exists either an indivisible good or 1% of a divisible good in player j 's bundle such that player i does not envy player j after removing it from player j 's bundle.

Final implementation:

Explicitly splitting each divisible good into 100 identical indivisible goods results in two computational challenges:

- (1) The number of goods, and, as a result, the number of decision variables in the resulting MILP increase significantly.
- (2) The number of constraints required to encode the piecewise-linear approximation of the logarithm function (in the form of segments or tangents on the log curve) is proportional to the number of possible utility levels that a player can achieve, which also increases from 1,000 to $1,000 \times 100$.

The former can be alleviated almost completely. Recall that the first step to computing the MNW solution is to find a largest set of players that can simultaneously derive a positive utility. This requires computing a maximum-cardinality matching, for which we use the MatlabBGL library.⁶ Since the maximum-cardinality algorithm works on sparse graphs and is extremely fast in practice, the increased number of goods is not an issue in this step.

The next step is to compute the MNW solution for the reduced set of players using the MILP of Figure 4. Here, the increased number of goods could affect the running time significantly. However, note that the indivisible goods created from a divisible good g are identical. Hence, we can retain the original decision variables $x_{i,g}$, but use them to denote the number of parts (out of 100) of good g that player i receives, rather than denoting whether player i receives good g entirely. In particular, for each divisible good g and each player i , we replace all the occurrences of $x_{i,g}$ in the

⁶https://www.cs.purdue.edu/homes/dgleich/packages/matlab_bgl.

MILP of Figure 4 with $x_{i,g}/100$, and replace $x_{i,g} \in \{0, 1\}$ with $x_{i,g} \in \{0, 1, \dots, 100\}$. The resulting MILP still has $n \cdot m$ integer (though, not binary) variables and n continuous variables, and we solve it using CPLEX.

Finally, for the latter challenge, note that although the number of possible utility levels that a player can achieve could, in the worst case, be 10^5 , in practice it is significantly smaller. We use a preprocessing step to identify the possible utility levels for each player using a variant of the standard dynamic programming algorithm for the Knapsack problem, implemented efficiently in MATLAB through vectorization.

B THE ELUSIVE COMBINATION OF EF1 AND PO

In this section, we provide examples of several candidate solutions that fail to achieve EF1 and PO together for additive valuations—two properties that are fairly easy to achieve individually. This serves as a backdrop to our argument that it is compelling—even surprising—that the MNW solution achieves the two properties together (Theorem 3.2).

Example B.1 (Rounding any MNW allocation for divisible goods violates EF1). The example we provide requires only 3 players but 31 goods. Let the set of players be $\mathcal{N} = \{1, 2, 3\}$. Suppose we have four types of goods: a single good of type a , and 10 goods each of types b , c , and d . Each player identically values all goods of the same type. Let the valuations of the players (specified only as a function of the type of the good) be as follows:

	Type a	Type b	Type c	Type d
Player 1	20	1	1.3	1.3
Player 2	15	0	1	1.3
Player 3	10	0	0	1

We show that when the goods are divisible, under the unique MNW allocation A^* , all goods of type b , c , and d are allocated entirely to players 1, 2, and 3, respectively, and the single good of type a is divided between players 1, 2, and 3 in fractions $10/18$, $7/18$, and $1/18$, respectively.

To verify this, we can use the KKT conditions. In allocation A^* , let U_1 , U_2 , and U_3 denote the utilities to the three players; thus, $U_1 = 10 + 20 \cdot (10/18)$, $U_2 = 10 + 15 \cdot (7/18)$, and $U_3 = 10 + 10 \cdot (1/18)$.

Now, the KKT condition at the good of type a is $20/U_1 = 15/U_2 = 10/U_3$, which is satisfied under A^* . The KKT condition at each good of type b is trivially satisfied under A^* , because no player other than player 1 likes goods of type b . At each good of type c , we need $1/U_2 \geq 1.3/U_1$, which is also satisfied under A^* . Finally, at each good of type d , we need $1/U_3 \geq \max\{1.3/U_1, 1.3/U_2\}$, which is also satisfied under A^* . Hence, A^* is the MNW allocation with divisible goods.

Let us now find an allocation for *indivisible* goods by rounding A^* . Because the allocation for divisible goods does not divide goods of types b , c , and d , no rounding scheme can alter the allocation of these goods. However, we now show that subject to this constraint, allocating the single good of type a entirely to any single player violates EF1. Indeed, if we allocate the good to player 1 (resp. player 2), player 2 (resp. player 1) envies player 3 even after removing any single good from player 3's bundle. If we allocate the good to player 3, player 1 envies player 2 even after removing any single good from player 2's bundle.

This shows that in this example, no rounding scheme applied to the unique MNW allocation for divisible goods can produce an EF1 allocation of indivisible goods. Because Theorem 3.2 asserts that an MNW allocation of indivisible goods is guaranteed to be EF1 and PO, this is also a fascinating example in which no way of rounding the MNW allocation for divisible goods produces an MNW allocation for indivisible goods. In other words, an MNW allocation for indivisible goods

inevitably gives at least one good to a player that receives a *zero* fraction of that good under the MNW solution for divisible goods.

Indeed, in this example, the unique MNW allocation for indivisible goods is as follows: Like the MNW allocation for divisible goods, it allocates all goods of types b and d to players 1 and 3, respectively. It allocates the good of type a to player 2, and to balance that, it allocates 9 out of the 10 goods of type c to player 1, and a single good of type c to player 2. Note that the 9 goods of type c that are now fully allocated to player 1 were fully allocated to player 2 in the MNW allocation for divisible goods.

In the economics literature, three popular notions of welfare—utilitarian, Nash, and egalitarian—are often arranged on a spectrum in which maximizing the utilitarian welfare is considered the most efficient, maximizing the egalitarian welfare is considered the fairest, and maximizing the Nash welfare is considered a good tradeoff between efficiency and fairness. In contrast, in our setting, maximizing Nash welfare is the only solution out of the three that achieves our desired fairness notion (EF1).

Example B.2 (Maximizing the utilitarian or the egalitarian welfare violates EF1). The fact that maximizing the utilitarian welfare violates EF1 is very easy to see. Let the set of players be $\mathcal{N} = \{1, 2\}$, the set of goods be $\mathcal{M} = \{g_1, g_2, g_3\}$, and the additive valuations of the players be as follows:

	g_1	g_2	g_3
Player 1	1/2	1/2	0
Player 2	2/5	2/5	1/5

Note that the unique allocation that maximizes the utilitarian welfare allocates goods g_1 and g_2 to player 1, and good g_3 to player 2, causing player 2 to envy player 1 even after removal of any single good from player 1's bundle.

To show that maximizing the egalitarian welfare violates EF1, we use a slightly more involved example. Let the set of players be $\mathcal{N} = \{1, 2, 3\}$, the set of goods be $\mathcal{M} = \{g_1, g_2, g_3, g_4\}$, and the additive valuations of the players be as follows:

	g_1	g_2	g_3	g_4
Player 1	1	0	0	0
Player 2	2/3	0	1/6	1/6
Player 3	0	1/5	2/5	2/5

First, to achieve a positive egalitarian welfare, we must allocate good 1 to player 1. Subject to this, the egalitarian welfare is uniquely maximized when good g_2 is allocated to player 3, and both goods g_3 and g_4 are allocated to player 2. However, this causes player 3 to envy player 2 even after removal of any single good from player 2's bundle.

Example B.3 (Maximizing the utilitarian/egalitarian welfare subject to EF1). The following counterexample shows that maximizing the utilitarian welfare subject to EF1 violates PO. This example was discovered using computer simulations. Let the set of players be $\mathcal{N} = \{1, 2, 3, 4\}$, the set of goods be $\mathcal{M} = \{g_i\}_{i \in [10]}$, and the additive valuations of the players be as follows:

	g_1	g_2	g_3	g_4	g_5	g_6	g_7	g_8	g_9	g_{10}
1	0.0426	0.0004	0.1019	0.1503	0.0541	0.1782	0.1212	0.0259	0.1574	0.1681
2	0.0365	0.0004	0.2311	0.1479	0.0649	0.1150	0.1501	0.1894	0.0285	0.0362
3	0.1124	0.0972	0.0574	0.0956	0.1441	0.1461	0.0674	0.1272	0.0254	0.1273
4	0.0368	0.0582	0.0242	0.0784	0.1844	0.1260	0.1124	0.1121	0.1610	0.1064

It can be checked that maximizing the utilitarian welfare subject to the EF1 constraint results in the following allocation A :

$$A_1 = \{g_6, g_7, g_{10}\}, A_2 = \{g_3, g_4, g_8\}, A_3 = \{g_1, g_2\}, \text{ and } A_4 = \{g_5, g_9\}.$$

However, this allocation is not Pareto optimal. An alternative allocation in which players 1 and 2 exchange goods g_7 and g_4 improves the utility to both players 1 and 2 while keeping the utility to both players 3 and 4 unaltered. This alternative allocation is not selected in the first place, because it violates EF1 (player 3 now envies player 1 even after the removal of any single good from player 1's bundle).

It is easy to see why maximizing the egalitarian welfare subject to EF1 violates PO. Suppose the set of players is $\mathcal{N} = \{1, 2, 3\}$, and the set of goods is $\mathcal{M} = \{g_1, g_2, g_3\}$. Let the valuations of the players be as follows:

	g_1	g_2	g_3
Player 1	2/3	1/3	0
Player 2	1/3	2/3	0
Player 3	1/3	1/3	1/3

Clearly the optimal egalitarian welfare is $1/3$ in this example. An EF1 allocation A that achieves this optimal welfare is given by $A_1 = \{g_2\}$, $A_2 = \{g_1\}$, and $A_3 = \{g_3\}$. However, this is clearly not PO: if players 1 and 2 exchange their bundles, they can both be better off without reducing the utility to player 3. Hence, maximizing the egalitarian welfare (at least naïvely) subject to EF1 is not PO.

It is worth remarking that in this example, replacing the egalitarian solution with its refinement, the leximin solution,⁷ produces an allocation that is EF1 and PO. While it is not clear if choosing the leximin-optimal allocation among the set of EF1 allocations satisfies PO, we conjecture that it does not. However, this will require a more complicated example.

While EF1 and PO are both mild properties by themselves, their combination is surprisingly elusive, which provides a justification for the MNW solution. However, EF1 and PO by themselves are not sufficient to always guarantee a desirable outcome. The following observations illustrate why.

Example B.4. Imagine we have a set of two players $\mathcal{N} = \{1, 2\}$, and a set of two goods $\mathcal{M} = \{g_1, g_2\}$. Suppose player 1 values both goods equally, and player 2 only values good g_2 .

In this case, the only intuitively fair outcome (which is also the outcome that the MNW solution selects) assigns good g_1 to player 1, and good g_2 to player 2. However, note that assigning both goods to player 1 also satisfies EF1 and PO, but is clearly undesirable.

More formally, we can argue that, while the MNW solution provides $\pi_n = 1/\Theta(\sqrt{n})$ -approximation of the MMS guarantee, simply restricting the allocation to be EF1 and PO gives a worse $1/n$ -approximation of MMS.

THEOREM B.5. *Every allocation that is envy-free up to one good (EF1) and Pareto optimal (PO) is $1/n$ -maximin share (MMS) for additive valuations over indivisible goods. Further, the factor $1/n$ is tight, i.e., for every $n \in \mathbb{N}$ and $\epsilon > 0$, there exists an instance with n players having additive valuations and an allocation satisfying EF1 and PO that is not $(1/n + \epsilon)$ -MMS.*

⁷The leximin solution maximizes the egalitarian welfare, but breaks ties among all allocations with the highest minimum utility in favor of those with a higher second minimum utility, and so on.

PROOF. We first prove that every allocation satisfying EF1 and PO is $1/n$ -MMS, and later prove tightness of the approximation ratio (upper bound).

Proof of the lower bound: Let A be an allocation satisfying EF1 and PO. Fix a player $i \in \mathcal{N}$. Because A is EF1, for every player $j \in \mathcal{N} \setminus \{i\}$ there exists a good $g_{ij} \in A_j$ such that

$$v_i(A_i) \geq v_i(A_j) - v_i(g_{ij}). \quad (6)$$

Let $T_i = \sum_{g \in \mathcal{M}} v_i(g)$, and let $E_i = \sum_{j \in \mathcal{N} \setminus \{i\}} v_i(g_{ij})$. Then, summing Equation (6) over all $j \in \mathcal{N} \setminus \{i\}$, we get

$$(n-1) \cdot v_i(A_i) \geq \sum_{j \in \mathcal{N} \setminus \{i\}} v_i(A_j) - E_i \Rightarrow n \cdot v_i(A_i) \geq T_i - E_i. \quad (7)$$

However, consider any partition of the set of goods $\mathcal{M}$ into n bundles. Due to the pigeonhole principle, there must exist a bundle that does not contain good g_{ij} for any $j \in \mathcal{N} \setminus \{i\}$. Since the value of this bundle according to player i can be at most $T_i - E_i$, the MMS guarantee of player i is also at most $T_i - E_i$. Equation (7) now implies that under A , each player i receives at least $1/n$ of her MMS guarantee, i.e., A is $1/n$ -MMS.

Proof of the upper bound: We now show that for every $n \in \mathbb{N}$ and $\epsilon > 0$, there exists an instance with n players for which some allocation satisfying EF1 and PO is not $(1/n + \epsilon)$ -MMS.

Construct an instance with n players and $2n-1$ goods. Let there be $n-1$ “high” goods that each player values at n , and n “low” goods that each player values at 1. The MMS guarantee of each player is n : the player can put each “high” good in its own bundle, and all “low” goods in a single bundle.

However, one can check that giving $n-1$ of the players a high and a low good each, and giving the remaining player the remaining single low good also satisfies EF1 and PO, but gives the last player exactly $1/n$ of her MMS guarantee. $\square$

C GENERAL VALUATIONS

In this section, we provide the definitions of the families of valuation functions mentioned in Section 3.1, and provide the missing proofs and examples. Let us begin by formally defining subadditive, superadditive, submodular, and supermodular valuations.

Definition C.1 (Subadditive and Superadditive Valuations). A valuation function $v : 2^{\mathcal{M}} \rightarrow \mathbb{R}_{\geq 0}$ is called subadditive (resp. superadditive) if for every pair of disjoint sets $S, T \subseteq \mathcal{M}$, we have $v(S \cup T) \leq v(S) + v(T)$ (resp. $v(S \cup T) \geq v(S) + v(T)$).

Definition C.2 (Submodular and Supermodular Valuations). A valuation function $v : 2^{\mathcal{M}} \rightarrow \mathbb{R}_{\geq 0}$ is called submodular (resp. supermodular) if for every pair $S, T \subseteq \mathcal{M}$, we have $v(S \cup T) \leq v(S) + v(T) - v(S \cap T)$ (resp. $v(S \cup T) \geq v(S) + v(T) - v(S \cap T)$).

It is clear that submodular (resp. supermodular) valuations are a special case of subadditive (resp. superadditive) valuations. We now provide a proof of Theorem 3.3, which asserts that for supermodular (and thus superadditive) valuations and subadditive valuations, EF1 and PO are incompatible.

PROOF OF THEOREM 3.3. Let the set of players be $\mathcal{N} = \{1, 2\}$, and the set of goods be $\mathcal{M} = \{a, b, c, d\}$. We use a common valuation for both players. Figures 7 and 6 define the supermodular (thus, superadditive) valuation v^{sup} and the subadditive valuation v^{sub} , respectively, through their value for a set $S \subseteq \mathcal{M}$.

In each case, under a PO allocation, player 1 receives one of the following sets of goods: $\emptyset$, $\{a\}$, $\{b, c, d\}$, and $\mathcal{M}$; and player 2 receives the set of remaining goods. It is easy to check that these are the only four PO allocation. Note that EF1 is violated in the first two allocations due to player 1

$$v^{\text{sub}}(S) = \begin{cases} 10 & \text{if } |S| = 4, \\ 7 & \text{if } |S| = 3, \\ 6 & \text{if } |S| = 2 \text{ and } a \notin S, \\ 4 & \text{if } |S| = 2 \text{ and } a \in S, \\ 4 & \text{if } S = \{a\}, \\ 3 & \text{if } S = \{b\}, \{c\}, \text{ or } \{d\}, \\ 0 & \text{if } S = \emptyset. \end{cases}$$

Fig. 6. Subadditive valuation.

$$v^{\text{sup}}(S) = \begin{cases} 4 & \text{if } |S| = 4, \\ 3 & \text{if } |S| = 3, \\ 2 & \text{if } |S| = 2 \text{ and } a \notin S, \\ 1 & \text{if } |S| = 2 \text{ and } a \in S, \\ 1 & \text{if } S = \{a\}, \\ 0 & \text{if } S = \{b\}, \{c\}, \{d\}, \text{ or } \emptyset. \end{cases}$$

Fig. 7. Supermodular (thus, superadditive) valuation.

envying player 2 (and in last two allocations due to player 2 envying player 1) even after removal of any single good from the envied player's bundle. $\square$

We now focus on the interesting case of submodular valuations, which are characteristic of substitute goods, and are alternatively defined via diminishing marginal utility. Examples of submodular valuations include unit demand valuations, strong valuations with no complementarities, and gross substitutes.

As mentioned in Section 3.1, we were unable to settle the question of the compatibility of EF1 and PO for submodular valuations. We know that an MNW allocation does not guarantee EF1 and PO for submodular valuations, but we can show that it guarantees MEF1 (a relaxation of EF1 that coincides with EF1 for additive valuations) and PO.

Example C.3 (MNW is not EF1 and PO for submodular valuations). Let the set of players be $\mathcal{N} = \{1, 2\}$, and the set of goods be $\mathcal{M} = \{a, b, c, d\}$. The submodular valuations v_1 and v_2 of players 1 and 2, respectively, are as follows:

Player 1: First, let us define the value of the player for individual goods.

$$v_1(a) = 1, v_1(b) = 1, v_1(c) = 0, \text{ and } v_1(d) = 0.$$

For $S \subseteq \mathcal{M}$ with $|S| \geq 2$, define $v_1(S)$ to be the sum of the values of the *two goods* in S that are the most valuable to player 1. It is easy to check that this is a submodular valuation.

Player 2: Let the value of the player for individual goods be as follows:

$$v_2(a) = 2.5, v_2(b) = 2.5, v_2(c) = 1, \text{ and } v_2(d) = 1.$$

Once again, for $S \subseteq \mathcal{M}$ with $|S| \geq 2$, define $v_2(S)$ to be the sum of the values of the *two goods* in S that are the most valuable to player 2. Similarly to v_1 , v_2 is also a submodular valuation.

Note that an MNW allocation must allocate at least one of the goods in $\{a, b\}$ to player 1 to achieve positive Nash welfare. If player 1 receives only one of these two goods, the Nash welfare can be at most $1 \cdot 3.5 = 3.5$. In contrast, the allocation A that gives both a and b to player 1 ($A_1 = \{a, b\}$), and the rest to player 2 ($A_2 = \{c, d\}$), achieves Nash welfare of $2 \cdot 2 = 4$. Hence, A is the unique MNW allocation.

However, this violates envy freeness up to one good (EF1). In particular, player 2 envies player 1 even after removal of any single good from player 1's bundle, because $v_2(A_2) = 2 < v_2(A_1 \setminus \{a\}) = v_2(A_1 \setminus \{b\}) = 2.5$.

In contrast, marginal envy freeness up to one good (MEF1) is satisfied, because $v_2(A_2) = 2 > v_2(A_2 \cup A_1 \setminus \{a\}) - v_2(A_2) = 3.5 - 2 = 1.5$.

We end this section with a proof of Theorem 3.5, which asserts that every MNW allocation is MEF1 and PO for submodular valuations.

PROOF OF THEOREM 3.5. Let $\mathbf{A}$ be an MNW allocation. First, let us prove the result for the case of $\text{NW}(\mathbf{A}) > 0$. In this case, the Pareto optimality of $\mathbf{A}$ is obvious due to the fact that $\mathbf{A}$ maximizes the Nash welfare. Suppose, for contradiction, that $\mathbf{A}$ is not MEF1. Then, there exist players $i, j \in N$ such that

$$\forall g \in A_j, v_i(A_i \cup A_j \setminus \{g\}) - v_i(A_i) > v_i(A_i). \quad (8)$$

Next, for every $r \in A_j$, let us define

$$\delta_i(g) = v_i(A_i \cup \{g\}) - v_i(A_i), \text{ and } \delta_j(g) = v_j(A_j) - v_j(A_j \setminus \{g\}).$$

Note that $\delta_i(g)$ and $\delta_j(g)$ are generalizations of $v_i(g)$ and $v_j(g)$ from additive valuations to submodular valuations. Also, observe that they are defined a bit differently for i and j .

We now derive two key results.

LEMMA C.4. For every $g^* \in A_j$, we have $\sum_{g \in A_j} \delta_i(g) > v_i(A_i \cup \{g^*\})$.

PROOF. Fix $g^* \in A_j$. Let us enumerate the elements of A_j as $g_1, \dots, g_k$ where $k = |A_j|$ and $g_k = g^*$. Also, for $t \in [k]$ define $A_j^t = \{g_1, \dots, g_t\}$, and $A_j^0 = \emptyset$. Then,

$$\begin{aligned} \sum_{g \in A_j \setminus \{g^*\}} \delta_i(g) &= \sum_{t=1}^{k-1} v_i(A_i \cup \{g_t\}) - v_i(A_i) \geq \sum_{t=1}^{k-1} v_i(A_i \cup A_j^t) - v_i(A_i \cup A_j^{t-1}) \\ &= v_i(A_i \cup A_j \setminus \{g^*\}) - v_i(A_i) > v_i(A_i), \end{aligned}$$

where the second transition holds due to submodularity of v_i , because the marginal value of adding g_t to A_i should be at least as much as the marginal value of adding g_t to $A_i \cup A_j^{t-1}$ (note that $A_j^t = A_j^{t-1} \cup \{g_t\}$). The final transition follows from Equation (8). Adding $\delta_i(g^*) = v_i(A_i \cup \{g^*\}) - v_i(A_i)$ on both sides yields the desired result. ■ (Proof of Lemma C.4)

LEMMA C.5. We have $\sum_{g \in A_j} \delta_j(g) \leq v_j(A_j)$.

PROOF. Once again, let $A_j = \{g_1, \dots, g_k\}$, where $k = |A_j|$, $A_j^t = \{g_1, \dots, g_t\}$ for $t \in [k]$, and $A_j^0 = \emptyset$. Then,

$$\sum_{g \in A_j} \delta_j(g) = \sum_{t=1}^k v_j(A_j) - v_j(A_j \setminus \{g_t\}) \leq \sum_{t=1}^k v_j(A_j^t) - v_j(A_j^{t-1}) = v_j(A_j),$$

where the inequality follows from the submodularity of v_j . ■ (Proof of Lemma C.5)

From Lemma C.4, it is clear that $\sum_{g \in A_j} \delta_i(g) > 0$. Thus, there exists $g \in A_j$ such that $\delta_i(g) > 0$. Fix $g^* = \arg \min_{g \in A_j: \delta_i(g) > 0} \delta_j(g) / \delta_i(g)$. We now take the ratio of the inequality in Lemma C.5 to the inequality in Lemma C.4 applied to our chosen g^* . This is well defined, because we already showed $\sum_{g \in A_j} \delta_i(g) > 0$, and we also have $v_i(A_i \cup \{g^*\}) \geq v_i(A_i) > 0$.

$$\frac{v_j(A_j)}{v_i(A_i \cup \{g^*\})} \geq \frac{\sum_{g \in A_j} \delta_j(g)}{\sum_{g \in A_j} \delta_i(g)} \geq \frac{\delta_j(g^*)}{\delta_i(g^*)} = \frac{v_j(A_j) - v_j(A_j \setminus \{g^*\})}{v_i(A_i \cup \{g^*\}) - v_i(A_i)},$$

where the second transition holds due to our choice of g^* . Upon rearranging the terms, we get

$$v_i(A_i \cup \{g^*\}) \cdot v_j(A_j \setminus \{g^*\}) > v_i(A_i) \cdot v_j(A_j),$$

which is a contradiction, because it implies that shifting g^* from player j to player i would increase the Nash welfare, which is in direct violation of the optimality of the Nash welfare under the MNW allocation $\mathbf{A}$.

Let us now handle the case of $\text{NW}(\mathbf{A}) = 0$. Let S denote the set of players that receive positive utility under $\mathbf{A}$. The proof of Pareto optimality of $\mathbf{A}$ for submodular valuations is identical to

the proof of Pareto optimality of an MNW allocation for additive valuations, which does not use additivity of the valuations. We now show that $\mathbf{A}$ is MEF1. Note that MEF1 holds among players in S due to the proof of the previous case, and holds trivially among players in $\mathcal{N} \setminus S$. Hence, the only case we need to address is when a player $i \in \mathcal{N} \setminus S$ (with $A_i = \emptyset$) marginally envies player $j \in S$ (with $v_j(A_j) > 0$) up to one good. Then, by the definition of MEF1, we have

$$\forall g \in A_j, v_i(A_j \setminus \{g\}) > 0. \quad (9)$$

Submodularity of v_j implies that $\sum_{g \in A_j} v_j(\{g\}) \geq v_j(A_j) > 0$. Hence, there exists a good $\hat{g} \in A_j$ such that $v_j(\{\hat{g}\}) > 0$. Applying Equation (9) to $\hat{g}$, we get $v_i(A_j \setminus \{\hat{g}\}) > 0$. But then moving all goods in A_j except $\hat{g}$ from player j to player i gives positive utility to player i while still giving positive utility to player j , which violates the fact that $\mathbf{A}$ provides positive utility to the maximum number of players. Hence, $\mathbf{A}$ must be MEF1. ■ (Proof of Theorem 3.5)

D A SPECTRUM OF FAIR AND EFFICIENT SOLUTIONS

In this article, we focused on the MNW solution that maximizes the Nash welfare, i.e., selects an allocation $\mathbf{A}$ that maximizes $\prod_{i \in \mathcal{N}} v_i(A_i)$. For simplicity, let us assume $\mathcal{N} = [n]$. Another popular solution concept for fair allocation, originally used in the networking literature, is *proportional fairness* [26].

Definition D.1 (Proportional Fairness). An allocation $\mathbf{A}$ is said to satisfy proportional fairness if for any alternative allocation $\mathbf{A}'$, it holds that

$$\sum_{i=1}^n \frac{v_i(A'_i) - v_i(A_i)}{v_i(A_i)} \leq 0.$$

In words, an allocation is proportionally fair if, when switching to an alternative allocation, the total percentage change in players' utilities is non-positive. This, in some sense, indicates that $\mathbf{A}$ is socially preferred to every other allocation.

The proportional fairness requirement can be equivalently written as

$$\sum_{i=1}^n \frac{v_i(A'_i)}{v_i(A_i)} \leq n \Leftrightarrow \frac{n}{\sum_{i=1}^n \left(\frac{v_i(A_i)}{v_i(A'_i)} \right)^{-1}} \geq 1.$$

That is, proportional fairness requires that the *harmonic mean* of the set of quantities $\{v_i(A_i)/v_i(A'_i)\}_{i \in [n]}$ be at least 1 for every other allocation $\mathbf{A}'$. Let us compare this to the requirement for an allocation $\mathbf{A}$ to be an MNW allocation: For every other allocation $\mathbf{A}'$, we must have

$$\prod_{i=1}^n v_i(A_i) \geq \prod_{i=1}^n v_i(A'_i) \Leftrightarrow \sqrt[n]{\prod_{i=1}^n \frac{v_i(A_i)}{v_i(A'_i)}} \geq 1.$$

That is, the MNW solution requires that the *geometric mean* of the same set of quantities $\{v_i(A_i)/v_i(A'_i)\}_{i \in [n]}$ be at least 1 for all other allocations $\mathbf{A}'$.

This inspires us to define a spectrum of properties for the allocation of indivisible goods where we require that the p th power mean of the same set of quantities be at least 1. Recall that the p th power mean of a set of non-negative numbers $\{x_i\}_{i \in [n]}$ is defined as $(\frac{1}{n} \sum_{i=1}^n x_i^p)^{1/p}$. The harmonic mean corresponds to $p = -1$, and the geometric mean corresponds to $p = 0$.

Definition D.2. For $p \in \mathbb{R}$, define $\Gamma(p)$ to be the set of allocations $\mathbf{A} \in \Pi_n(\mathcal{M})$ such that for every other allocation $\mathbf{A}' \in \Pi_n(\mathcal{M})$, we have

$$\left(\frac{1}{n} \sum_{i=1}^n \left(\frac{v_i(A_i)}{v_i(A'_i)} \right)^p \right)^{1/p} \geq 1.$$

D.1 Relations Among $\Gamma(p)$

We now observe relations between the sets $\Gamma(p)$ for different values of p . First, the power-mean inequality states that for $p > p'$, the p th power mean is no less than the p' th power mean. This directly yields that $\Gamma(p)$ becomes stricter as p decreases.

THEOREM D.3 (DECREASING POWER). For $p, p' \in \mathbb{R}$ with $p' > p$, $\Gamma(p) \subseteq \Gamma(p')$.

The limiting cases of $p \rightarrow -\infty$ and $p \rightarrow \infty$ are interesting. Due to Theorem D.3, we say that an allocation $\mathbf{A}$ belongs to $\Gamma(-\infty)$ if it belongs to $\Gamma(p)$ for all $p \in \mathbb{R}$, and say that $\mathbf{A}$ belongs to $\Gamma(+\infty)$ if it belongs to $\Gamma(p)$ for some $p \in \mathbb{R}$. It is easy to observe that $\mathbf{A} \in \Gamma(-\infty)$ if it weakly Pareto dominates every other allocation $\mathbf{A}'$, i.e., if $v_i(A_i) \geq v_i(A'_i)$ for all $i \in N$ and $\mathbf{A}'$. This is an extremely stringent requirement. However, an allocation $\mathbf{A} \in \Gamma(+\infty)$ if it is Pareto optimal, which is a much weaker requirement. Consequently, we have the following result:

THEOREM D.4 (EFFICIENCY). For every $p \in \mathbb{R}$, every allocation $\mathbf{A} \in \Gamma(p)$ is Pareto optimal (PO).

PROOF. Indeed, assume that an allocation $\mathbf{A} \in \Gamma(p)$ is not PO. Let $\mathbf{A}' \in \Pi_n(\mathcal{M})$ be an allocation that Pareto dominates $\mathbf{A}$, i.e., $v_i(A'_i) \geq v_i(A_i)$ for all $i \in N$, and $v_{i^*}(A'_{i^*}) > v_{i^*}(A_{i^*})$ for some $i^* \in N$. Then, we would have

$$\sum_{i=1}^n \left(\frac{v_i(A_i)}{v_i(A'_i)} \right)^p < n,$$

which is a contradiction, because $\mathbf{A} \in \Gamma(p)$. $\square$

D.2 $\Gamma(0)$ is Special

$\Gamma(0)$, which is the set of MNW allocations, not only lies at the center of the spectrum, it is also special in two other ways. First, it is the strictest set that is guaranteed to be non-empty. We know that there trivially exists an MNW allocation, which belongs to $\Gamma(0)$. In contrast, there exist instances in which $\Gamma(p)$ is empty for every $p < 0$; simply consider a single good and two players having value 1 for the good.

THEOREM D.5 (EXISTENCE). For $p \in \mathbb{R}$, $\Gamma(p) \neq \emptyset$ if and only if $p \geq 0$.

Second, in a given instance, even if $\Gamma(p) \neq \emptyset$ for $p < 0$, this does not offer a refinement of $\Gamma(0)$, as the following result shows. For an allocation $\mathbf{A}$, we refer to the vector of player utilities $(v_i(A_i))_{i \in N}$ as the utility vector of $\mathbf{A}$. Note that the definition of $\Gamma(p)$ only relies on the utility vector of an allocation; thus, two allocations with identical utility vectors are essentially equivalent.

THEOREM D.6 (NO REFINEMENT). If $\Gamma(0)$ contains two allocations with distinct utility vectors, then $\Gamma(p) = \emptyset$ for all $p < 0$.

PROOF. Let $\mathbf{A}$ and $\mathbf{A}'$ be two arbitrary allocations in $\Gamma(0)$ with distinct utility vectors. Then, $\prod_{i \in N} v_i(A_i) = \prod_{i \in N} v_i(A'_i)$, implying that the geometric mean of $X_{\mathbf{A}, \mathbf{A}'} = \{v_i(A_i)/v_i(A'_i)\}_{i \in N}$ as well as the geometric mean of $X_{\mathbf{A}', \mathbf{A}} = \{v_i(A'_i)/v_i(A_i)\}_{i \in N}$ are both 1. Because $\mathbf{A}$ and $\mathbf{A}'$ have distinct utility vectors in $\Gamma(0)$, we know that not all the numbers in $X_{\mathbf{A}, \mathbf{A}'}$ (resp. $X_{\mathbf{A}', \mathbf{A}}$) are identical. Then, due to power-mean inequality, we have that for any $p < 0$, the p th power mean of $X_{\mathbf{A}, \mathbf{A}'}$ and

the p th power mean of $X_{A',A}$ are both strictly less than 1, implying that neither A nor A' belongs to $\Gamma(p)$. Because we chose A and A' to be arbitrary, it follows that $\Gamma(p) = \emptyset$, as required. $\square$

In other words, for $p < 0$, $\Gamma(p)$ can be non-empty only if all MNW allocations yield identical utility vectors, and $\Gamma(p) = \Gamma(0)$. This has several interesting consequences. First, we know that $\Gamma(0)$ is the strictest refinement that the spectrum offers. Second, if there exists a proportionally fair allocation (in the sense of Definition D.1), then an allocation is proportionally fair if and only if it is an MNW allocation.

D.3 Fairness Properties of the Spectrum

While Theorem D.4 establishes an allocation in $\Gamma(p)$ (for any value of p) is always Pareto optimal, it is unclear if an allocation in $\Gamma(p)$ is always fair. We already know that an MNW allocation (i.e., an allocation in $\Gamma(0)$) is EF1 (Theorem 3.2). A folklore result (for which we provide a simple proof below) states that a proportionally fair allocation (i.e., an allocation in $\Gamma(-1)$) is EF. We somewhat extend this fairness guarantee to the positive side of the spectrum ($p > 0$) by showing that an allocation in $\Gamma(1)$ is “EF2.”

Let us first introduce a family of relaxations of envy freeness : We say that an allocation A is envy-free up to k goods (EF k) if

$$\forall i, j \in \mathcal{N}, \exists S \subseteq A_j \text{ with } |S| \leq k, v_i(A_i) \geq v_i(A_j \setminus S).$$

THEOREM D.7 (FAIRNESS). *For $p \in [-1, 1]$, every allocation in $\Gamma(p)$ is envy-free up to $1 + \lceil p \rceil$ goods, where $\lceil \cdot \rceil$ is the ceiling function.*

PROOF. Due to Theorem D.3, we only need to prove this theorem for $p \in \{-1, 0, 1\}$. For $p = 0$, we already showed that every MNW allocation is EF1 (Theorem 3.2).

Let us now consider $p = -1$. Let allocation $A \in \Gamma(-1)$. Consider a pair of players j, j' . For every good $t \in A_{j'}$, we apply the inequality in the definition of $\Gamma(-1)$ using the allocation A and the allocation A' obtained by moving good t from player j' to player j . We have

$$\begin{aligned} n &\geq \sum_{i=1}^n \frac{v_i(A'_i)}{v_i(A_i)} = \left(\sum_{i \neq j, j'} \frac{v_i(A_i)}{v_i(A_i)} \right) + \frac{v_j(A_j \cup \{t\})}{v_j(A_j)} + \frac{v_{j'}(A_{j'} \setminus \{t\})}{v_{j'}(A_{j'})} \\ &= n + \frac{v_j(t)}{v_j(A_j)} - \frac{v_{j'}(t)}{v_{j'}(A_{j'})}, \end{aligned}$$

which implies that $v_j(t) \leq v_{j'}(t) \cdot v_j(A_j) / v_{j'}(A_{j'})$. Summing over all $t \in A_{j'}$, we get

$$v_j(A_{j'}) = \sum_{t \in A_{j'}} v_j(t) \leq \sum_{t \in A_{j'}} v_{j'}(t) \frac{v_j(A_j)}{v_{j'}(A_{j'})} = v_j(A_j),$$

i.e., player j is not envious for player j' .

Let us now consider $p = 1$. Consider an allocation $A \in \Gamma(1)$. Consider players i and j . We want to show that player i would not envy j if we are allowed to remove (up to) two goods from player j 's bundle. If $|A_j| \leq 2$, we are done. Assume $|A_j| \geq 3$. We now show that there are goods $t_1, t_2 \in A_j$ such that $v_i(A_i) \geq v_i(A_j \setminus \{t_1, t_2\})$.

Consider a good $t \in A_j$, and define the allocation A' obtained by moving good t from player j to player i in A . By the definition of $\Gamma(1)$, we get

$$\frac{v_i(A_i)}{v_i(A_i) + v_i(t)} + \frac{v_j(A_j)}{v_j(A_j) - v_j(t)} \geq 2.$$

Setting $x_t = \frac{v_i(t)}{v_i(A_i)}$ and $y_t = \frac{v_j(t)}{v_j(A_j)}$, this inequality can be expressed as

$$\frac{1}{1+x_t} + \frac{1}{1-y_t} \geq 2,$$

which implies that

$$x_t \leq \frac{y_t}{1-2y_t} \quad (10)$$

whenever $y_t \leq 1/3$. Now let t_1 and t_2 be the goods in A_j for which player j has the highest and second-highest value, respectively. Hence, for every good $t \in A_j \setminus \{t_1, t_2\}$, we have $y_t \leq 1/3$. Using Equation (10), we obtain that the value of player i for the goods in $A_j \setminus \{t_1, t_2\}$ is

$$\begin{aligned} v_i(A_j \setminus \{t_1, t_2\}) &= v_i(A_i) \sum_{t \in A_j \setminus \{t_1, t_2\}} \frac{v_i(t)}{v_i(A_i)} = v_i(A_i) \sum_{t \in A_j \setminus \{t_1, t_2\}} x_t \\ &\leq v_i(A_i) \sum_{t \in A_j \setminus \{t_1, t_2\}} \frac{y_t}{1-2y_t} \leq v_i(A_i) \sum_{t \in A_j \setminus \{t_1, t_2\}} \frac{y_t}{1-2y_{t_2}} \\ &\leq v_i(A_i), \end{aligned}$$

where the second inequality holds, because $y_t \leq y_{t_2}$ for $t \in A_j \setminus \{t_1, t_2\}$, and the final transition follows due to the definitions of y , t_1 , and t_2 . We thus have that **A** is EF2. $\square$

D.4 Computational Properties of the Spectrum

We proved that an allocation in $\Gamma(1)$ provides a slightly weaker fairness guarantee (EF2) than an MNW allocation in $\Gamma(0)$ does (EF1). Because computing an MNW allocation is $\mathcal{NP}$ -hard [33], one may wonder if an allocation in $\Gamma(1)$ can be computed in polynomial time, thus offering a computational advantage over the MNW solution at the expense of the fairness guarantee. Interestingly, a polynomial-time Turing reduction from the popular $\mathcal{NP}$ -hard PARTITION problem shows that computing an allocation in $\Gamma(p)$ is $\mathcal{NP}$ -hard for $p \in (0, 1]$. Note that it is the search problem (of actually finding the allocation) that is $\mathcal{NP}$ -hard rather than the decision problem of determining the existence of such an allocation (which is a trivial problem, as such an allocation always exists).

THEOREM D.8 (COMPUTATIONAL HARDNESS). *For $p \in [0, 1]$, computing an allocation in $\Gamma(p)$ is $\mathcal{NP}$ -hard.*

PROOF. Due to Theorem D.3, we only need to show the hardness for $p = 1$. We show that a polynomial-time algorithm to compute an allocation in $\Gamma(1)$ can be used to decide the PARTITION problem in polynomial time. The input in an instance of the PARTITION problem is a set of m positive integers $S = \{x_1, \dots, x_m\}$, and our goal is to decide whether there exists a *perfect partition* of S , i.e., a partition of S into two exclusive and exhaustive subsets whose sum of elements is equal. Let $T = \sum_{i \in [m]} x_i$. We say that a partition of S is a *minimum-difference partition* if the difference between the sums of the two subsets is the least possible among all partitions of S into two subsets.

Let us first construct a new set of $m' = m + 2$ positive integers $S' = \{x'_i\}_{i \in [m']}$ where $x'_i = 5x_i$ for $i \in [m]$, $x_{m+1} = 1$, and $x_{m+2} = 2$. Let $T' = \sum_{i \in [m']} x'_i = 5T + 3$. Note that S' does not have a perfect partition. Further, it has a minimum-difference partition with difference 1 if and only if S has a perfect partition. Note that a partition of S' with difference 1 can only be created by taking a perfect partition of S , replacing the elements of S by the corresponding elements of S' , and then adding x'_{m+1} and x'_{m+2} in different subsets.

Next, we construct an instance of our fair allocation problem as follows: We have two players with the identical valuation v over the set of goods $\mathcal{M} = [m']$ under which $v(i) = x'_i$ for each good

$i \in \mathcal{M}$. We can interpret an allocation $\mathbf{A}$ of this instance as a partition of S' , in which each subset is formed by taking the elements of S' corresponding to the goods in a player's bundle. Thus, the sums of the two subsets in the partition are exactly $v(A_1)$ and $v(A_2)$, and $v(A_1) + v(A_2) = T'$.

We now show that every allocation in $\Gamma(1)$ produces a minimum-difference partition of S' . To see this, consider an allocation $\mathbf{A}$ in $\Gamma(1)$, and without loss of generality, assume $v(A_1) - v(A_2) = \delta > 0$. Thus, $v(A_1) = (T' + \delta)/2$ and $v(A_2) = (T' - \delta)/2$. Now, suppose for contradiction that there exists another allocation $\mathbf{A}'$ under which $|v(A'_1) - v(A'_2)| = \epsilon < \delta$. Because S' does not admit a perfect partition, we have $\epsilon > 0$. Without loss of generality, let $v(A'_1) - v(A'_2) = \epsilon$ (otherwise, we can switch the bundles of the two players). Hence, $v(A'_1) = (T' + \epsilon)/2$ and $v(A'_2) = (T' - \epsilon)/2$. However, in this case

$$\frac{v(A_1)}{v(A'_1)} + \frac{v(A_2)}{v(A'_2)} - 2 = \frac{T' + \delta}{T' + \epsilon} + \frac{T' - \delta}{T' - \epsilon} - 2 = (\delta - \epsilon) \cdot \left[\frac{1}{T' + \epsilon} - \frac{1}{T' - \epsilon} \right] < 0,$$

which contradicts the fact that $\mathbf{A}$ is an allocation in $\Gamma(1)$. Hence, $\mathbf{A}$ must produce a minimum-difference partition of S' .

To solve the original PARTITION instance, we compute an allocation in $\Gamma(1)$, use it to produce a minimum-difference partition of S' , and check if its difference is 1. $\square$

Thus, proportional fairness and the MNW solution, which coincide for allocation of divisible goods, are connected on a spectrum in the case of indivisible goods. The spectrum allows us to view the MNW solution as a strictest refinement of the set of Pareto optimal allocations that is guaranteed to be non-empty. Further, weaker solutions on the spectrum that exhibit weaker fairness guarantee do not offer a computational advantage over the MNW solution.

ACKNOWLEDGMENTS

We sincerely thank the anonymous referees of this article and of its conference version for their very helpful feedback.

REFERENCES

- [1] M. Aleksandrov, H. Aziz, S. Gaspers, and T. Walsh. 2015. Online fair division: Analysing a food bank problem. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI'15)*. 2540–2546.
- [2] G. Amanatidis, E. Markakis, A. Nikzad, and A. Saberi. 2015. Approximation algorithms for computing maximin share allocations. In *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP'15)*. 39–51.
- [3] N. Anari, S. O. Gharan, T. Mai, and V. V. Vazirani. 2018. Nash social welfare for indivisible items under separable, piecewise-linear concave utilities. In *Proceedings of the 29th ACM-SIAM Symposium on Discrete Algorithms (SODA'18)*.
- [4] N. Anari, S. O. Gharan, A. Saberi, and M. Singh. 2017. Nash social welfare, matrix permanent, and stable polynomials. In *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS'17)*.
- [5] K. J. Arrow and M. D. Intriligator (Eds.). 1982. *Handbook of Mathematical Economics*. North-Holland.
- [6] H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. 2015. Fair assignment of indivisible objects under ordinal preferences. *Artific. Intell.* 227 (2015), 71–92.
- [7] J. Bartholdi, C. A. Tovey, and M. A. Trick. 1989. The computational difficulty of manipulating an election. *Soc. Choice. Welf.* 6 (1989), 227–241.
- [8] M. Berliant, K. Dunz, and W. Thomson. 1992. On the fair division of a heterogeneous commodity. *J. Math. Econ.* 21 (1992), 201–216.
- [9] A. Bogomolnaia and H. Moulin. 2001. A new solution to the random assignment problem. *J. Econ. Theor.* 100 (2001), 295–328.
- [10] A. Bogomolnaia and H. Moulin. 2004. Random matching under dichotomous preferences. *Econometrica* 72 (2004), 257–279.
- [11] S. Bouveret, U. Endriss, and J. Lang. 2010. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI'10)*. 387–392.
- [12] S. Bouveret and M. Lemaître. 2016. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonom. Agents Multi-Agent Syst.* 30, 2 (2016), 259–290.

- [13] S. J. Brams, M. Kilgour, and C. Klamler. 2017. Maximin envy-free division of indivisible items. *Group Decision and Negotiation* 26, 1 (2017), 115–131.
- [14] S. J. Brams and D. L. King. 2005. Efficient fair division: Help the worst off or avoid envy? *Ration. Soc.* 17, 4 (2005), 387–421.
- [15] S. J. Brams and A. D. Taylor. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.
- [16] E. Budish. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *J. Polit. Econ.* 119, 6 (2011), 1061–1103.
- [17] E. Budish, Y.-K. Che, F. Kojima, and P. Milgrom. 2013. Designing random allocation mechanisms: Theory and applications. *Amer. Econ. Rev.* 103, 2 (2013), 585–623.
- [18] R. Cole, N. Devanur, V. Gkatzelis, K. Jain, T. Mai, V. V. Vazirani, and S. Yazdanbod. 2017. Convex program duality, Fisher markets, and Nash social welfare. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC'17)*. 459–460.
- [19] R. Cole and V. Gkatzelis. 2015. Approximating the Nash social welfare with indivisible items. In *Proceedings of the 47th ACM Symposium on Theory of Computing (STOC'15)*. 371–380.
- [20] V. Conitzer, R. Freeman, and N. Shah. 2017. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC'17)*. 629–646.
- [21] B. de Keijzer, S. Bouveret, T. Klos, and Y. Zhang. 2009. On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences. In *Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT'09)*. 98–110.
- [22] E. Eisenberg and D. Gale. 1959. Consensus of subjective probabilities: The pari-mutuel method. *Ann. Math. Stat.* 30, 1 (1959), 165–168.
- [23] D. Foley. 1967. Resource allocation and the public sector. *Yale Economics Essays* 7 (1967), 45–98.
- [24] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. 2011. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation (NSDI'11)*. 24–37.
- [25] J. Goldman and A. D. Procaccia. 2014. Spliddit: Unleashing fair division algorithms. *SIJecom Exch.* 13, 2 (2014), 41–46.
- [26] F. P. Kelly. 1997. Charging and rate control for elastic traffic. *Euro. Trans. Telecomm.* 8 (1997), 33–37.
- [27] D. Kurokawa, A. D. Procaccia, and N. Shah. 2015. Leximin allocations in the real world. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC'15)*. 345–362.
- [28] D. Kurokawa, A. D. Procaccia, and J. Wang. 2016. Can the maximin share guarantee be guaranteed? In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI'16)*. 523–529.
- [29] E. Lee. 2015. APX-Hardness of Maximizing Nash Social Welfare with Indivisible Items. Retrieved from: [arXiv:1507.01159](https://arxiv.org/abs/1507.01159).
- [30] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 6th ACM Conference on Economics and Computation (EC'04)*. 125–131.
- [31] H. Moulin. 2003. *Fair Division and Collective Welfare*. The MIT Press, Cambridge, MA.
- [32] J. Nash. 1950. The bargaining problem. *Econometrica* 18, 2 (1950), 155–162.
- [33] T. T. Nguyen, M. Roos, and J. Rothe. 2013. A survey of approximability and inapproximability results for social welfare optimization in multiagent resource allocation. *Ann. Math. Artif. Intell.* 68, 1 (2013), 65–90.
- [34] A. Nongillard, P. Mathieu, and B. Jaumard. 2009. A realistic approach to solve the Nash welfare. In *Proceedings of the 7th International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS'09)*. 374–382.
- [35] E. Pazner and D. Schmeidler. 1978. Egalitarian equivalent allocations: A new concept of economic equity. *Quart. J. Econ.* 92, 4 (1978), 671–687.
- [36] A. D. Procaccia and J. Wang. 2014. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the 14th ACM Conference on Economics and Computation (EC'14)*. 675–692.
- [37] S. Ramezani and U. Endriss. 2010. Nash social welfare in multiagent resource allocation. In *Proceedings of the 12th International Workshop on Agent-Mediated Electronic Commerce (AMEC'10)*. 117–131.
- [38] J. Schummer. 1997. Strategy-proofness versus efficiency on restricted domains of exchange economies. *Soc. Choice Welf.* 14 (1997), 47–56.
- [39] H. Steinhilber. 1948. The problem of fair division. *Econometrica* 16 (1948), 101–104.
- [40] B. Sziklai and E. Segal-Halevi. 2015. Resource-monotonicity and Population-monotonicity in Cake-cutting. Retrieved from: [arXiv:1510.05229](https://arxiv.org/abs/1510.05229).
- [41] H. Varian. 1974. Equity, envy and efficiency. *J. Econ. Theor.* 9 (1974), 63–91.
- [42] D. Weller. 1985. Fair division of a measurable space. *J. Math. Econ.* 14, 1 (1985), 5–17.
- [43] E. P. Wigner. 1960. The unreasonable effectiveness of mathematics in the natural sciences. *Commun. Pure Appl. Math.* 13, 1 (1960), 1–14.

Received January 2017; revised October 2017; accepted June 2018

ACM Transactions on Economics and Computation, Vol. 7, No. 3, Article 12. Publication date: September 2019.