



Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A. D., Shah, N., and Wang, J. (2016) The Unreasonable Fairness of Maximum Nash Welfare. In: EC '16 17th ACM Conference on Economics and Computation, Maastricht, The Netherlands, 24-28 Jul 2016, pp. 305-322. ISBN 9781450339360.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/123283/>

Deposited on: 24 August 2016

Enlighten – Research publications by members of the University of Glasgow
<http://eprints.gla.ac.uk>

The Unreasonable Fairness of Maximum Nash Welfare

IOANNIS CARAGIANNIS, University of Patras

DAVID KUROKAWA, Carnegie Mellon University

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

ARIEL D. PROCACCIA, Carnegie Mellon University

NISARG SHAH, Carnegie Mellon University

JUNXING WANG, Carnegie Mellon University

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 unexpectedly, 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 lead us to believe that MNW is the ultimate 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

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 $v_i(x_i) = \sum_g x_{ig}v_i(g)$, 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 [1974]. 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

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”; and by the Caratheodory research grant E.114 from the University of Patras.

Authors’ addresses: I. Caragiannis, University of Patras, Greece; email: caragian@ceid.upatras.gr; D. Kurokawa, Carnegie Mellon University, USA; email: dkurokaw@cs.cmu.edu; H. Moulin, University of Glasgow, UK, Higher School of Economics, St Petersburg, Russia; email: Herve.Moulin@glasgow.ac.uk; A. D. Procaccia, Carnegie Mellon University, USA; email: arielpro@cs.cmu.edu; N. Shah, Carnegie Mellon University, USA; email: nkshah@cs.cmu.edu; J. Wang, Carnegie Mellon University, USA; email: junxingw@cs.cmu.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 ACM 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.

EC’16, July 24–28, 2016, Maastricht, The Netherlands.

ACM 978-1-4503-3936-0/16/07 ...\$15.00.

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

<http://dx.doi.org/10.1145/http://dx.doi.org/10.1145/2940716.2940726>

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* [Foley 1967]: 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 chose 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 x that maximize the *Nash social welfare* $\prod_i v_i(x_i)$ [Arrow and Intriligator 1982, Volume 2, Chapter 14]. Consequently, a CEEI allocation can be computed in polynomial time via the convex program of Eisenberg and Gale [1959].

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 [Ramezani and Endriss 2010; Cole and Gkatzelis 2015]. 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 [Moulin 2003]. 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 paper is to demonstrate the “unreasonable effectiveness” [Wigner 1960] — 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 [Budish 2011; Ghodsi et al. 2011; Aleksandrov et al. 2015; Procaccia and Wang 2014; Kurokawa et al. 2015]. We are especially excited about making a real-world impact through *Spliddit* (www.spliddit.org), a not-for-profit fair division website [Goldman and Procaccia 2014]. Since its launch in November 2014, the website has attracted more than 60,000 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 current solution leaves something to be desired; here is how it works. First, to express their preferences, users simply need to divide 1000 points between the goods. This simple elicitation process relies on the additivity assumption, and is the reason why, in our view, it 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 [Procaccia and Wang 2014]. This allows Spliddit to provide a *provable* fairness guarantee for indivisible goods. That said, a (full) maximin share allocation can always be found in practice [Bouveret and Lemaître 2016; Kurokawa et al. 2016].

While the algorithm generally provides good solutions, it is highly discontinuous, and its direct reliance on the maximin share alone — when envy freeness and proportionality cannot be obtained — sometimes leads to nonintuitive outcomes. For example, consider this excerpt from an email from 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 seeks 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 incumbent algorithm for allocating goods (and to every other approach we know of, as we discuss below). It has been deployed on Spliddit on 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). 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 [Procaccia and Wang 2014]. 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 1281 real-world fair division instances from 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 [Nguyen et al. 2013]. 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 1000 points among the available goods. Our algorithm scales very well, solving relatively large instances with 50 players and 150 goods in less than 30 seconds, 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. [2004]. 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. [2004] does not guarantee additional properties.

Budish [2011] 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.

From an algorithmic perspective, Ramezani and Endriss [2010] show that maximizing Nash welfare is \mathcal{NP} -hard under certain combinatorial bidding languages (including, under additive valuations). Cole and Gkatzelis [2015] give a constant-factor, polynomial-time approximation under additive valuations (to be precise, their objective function is the geometric mean of the utilities).¹ Lee [2015] shows that the problem is APX-hard, that is, a constant-factor approximation is the best one can hope for.

When there are only two players, compelling approaches for allocating goods are available. In fact, Spliddit currently handles this case separately, via the Adjusted Winner algorithm [Brams and Taylor 1996]. The shortcoming of Adjusted Winner is that it usually 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 [1978], which is defined for any number of players. For $n > 2$ players, it may need to split all the goods, that is, 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 [Brams et al. 2015; Aziz et al. 2015]. This assumption (arguably)

¹However, a constant-factor approximation need not satisfy any of the theoretical guarantees we establish in this paper for the MNW solution.

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 [Bogomolnaia and Moulin 2001, 2004; Budish et al. 2013]; this is hardly appropriate for the cases we find on Spliddit in which the outcome is used only once.

Finally, it is worth noting that the idea of maximizing the product of utilities was studied by Nash [1950], 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 [Kelly 1997]: 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 paper, 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 paper 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 [Bouveret and Lemaître 2016; Procaccia and Wang 2014]. 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 [Brams and Taylor 1996] 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 [Bouveret and Lang 2008]. To guarantee existence, a somewhat weaker definition is called for; the following definition is a rather minimal relaxation.

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) if²

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

²To be perfectly accurate, this is not satisfied if A_j is empty, but, clearly, in this case i does not envy j .

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 paper, 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* [Budish 2011]. 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

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

We say that 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 [Procaccia and Wang 2014; Kurokawa et al. 2016], a $(2/3 + O(1/n))$ -MMS allocation always exists [Procaccia and Wang 2014], and can be computed in polynomial time [Amanatidis et al. 2015]. 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.³ To this end, we use the rather unrestrictive notion of *Pareto optimality*.

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

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

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.

³In the absence of this requirement, even envy freeness can be achieved by simply not allocating any goods.

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 (example presented in the full version of the paper⁴). Interestingly, maximizing these objectives subject to the constraint that the allocation is EF1 violates PO (example presented in the full version).

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 [Varian 1974], 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 an example in the full version.

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

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

If it is possible to achieve positive Nash welfare (i.e., provide positive utility to every player simultaneously), any Nash-welfare-maximizing allocation can be selected. In the special case that every feasible allocation has zero Nash welfare (i.e., it is impossible to provide positive utility to every player simultaneously), we find a largest set of players to which we can simultaneously provide positive utility, and select an allocation to these players maximizing their product of utilities. While this edge case is highly unlikely to appear in practice, it must be handled carefully to retain the solution’s attractive fairness and efficiency properties. We say that an allocation is a *maximum Nash welfare (MNW)* allocation if it can be selected by the MNW solution. A formal specification of the MNW solution is presented in the full version.

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 A denote an MNW allocation. First, let us assume $NW(A) > 0$. Pareto optimality of 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 A . Suppose, for contradiction, that 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 A' denote the allocation obtained by moving g^* from player j to player i in A . We now show that $NW(A') > NW(A)$, which gives the desired contradiction as the Nash

⁴Available at: <http://www.cs.cmu.edu/~arielp/papers.html>

welfare is optimal under A . Specifically, we show that $\text{NW}(A')/\text{NW}(A) > 1$. The ratio is well-defined because we assumed $\text{NW}(A) > 0$.

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,

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

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 $\text{NW}(A) = 0$. Let S denote the set of players to which the solution gives positive utility. Then, by the definition of the MNW solution, S is a largest set of players to which one can provide positive utility. Pareto optimality of 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 A already maximizes.

From the proof of the case of $\text{NW}(A) > 0$, we already know that there is no envy up to one good among players in S because A is an MNW allocation over these players, and under 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 A is both EF1 and PO. ■

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. We ask whether the same guarantees can be achieved for 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 the full version. 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 (an example is presented

in the full version) 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 $A \in \Pi_n(\mathcal{M})$ satisfies MEF1 if

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

Note that MEF1 is strictly weaker than EF1. However, 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 $2/3 + O(1/n)$ [Procaccia and Wang 2014], 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 [Procaccia and Wang 2014; Amanatidis et al. 2015]. 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 A be an MNW allocation. As in the proof of Theorem 3.2, we begin by assuming $\text{NW}(A) > 0$, and handle the case of $\text{NW}(A) = 0$ 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 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 $f(x; \beta)$ denote the RHS of the inequality above. Then, we can write $\beta \leq f(x; \beta) \leq \max_x f(x; \beta)$. 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 $f(x; \beta)$ over all 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, \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 which 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).

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

Player i , for $i \geq 2$: $v_i(h_i) = \frac{1}{\pi_n + 1}$, $v_i(l_i) = \frac{\pi_n}{\pi_n + 1}$, 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 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 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, $\text{NW}(A^*) > 0$, 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 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 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 A^* . Hence, the maximum Nash welfare of 1 is uniquely achieved by the allocation 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 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 on 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 $A \in \Pi_n(\mathcal{M})$ is an α -pairwise maximin share (MMS) allocation if

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

We simply say that 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 neither the pairwise MMS guarantee nor the MMS guarantee imply the other, it can be shown that a pairwise MMS allocation is $(1/2)$ -MMS (refer to the full version for a proof).

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 (refer to the full version for a proof). This, in particular, implies that pairwise MMS is stronger than EF1.

Definition 4.4 (EFX: Envy freeness up to the Least Valued Good). We say that an allocation $A \in \Pi_n(\mathcal{M})$ is envy free up to the least (positively) valued good if

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

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.

Given this motivation 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. The proof of the result appears in the full version.

COROLLARY 4.5. *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.*

4.3. Approximate MMS and Pairwise MMS, in Practice

Theorem 4.1 and Corollary 4.5 show that the MNW solution is *guaranteed* to be π_n -MMS and Φ -pairwise MMS. We now evaluate it on these benchmarks (which, we reiterate, it is *not* designed to optimize) using real-world data. Specifically, we use 1281 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 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

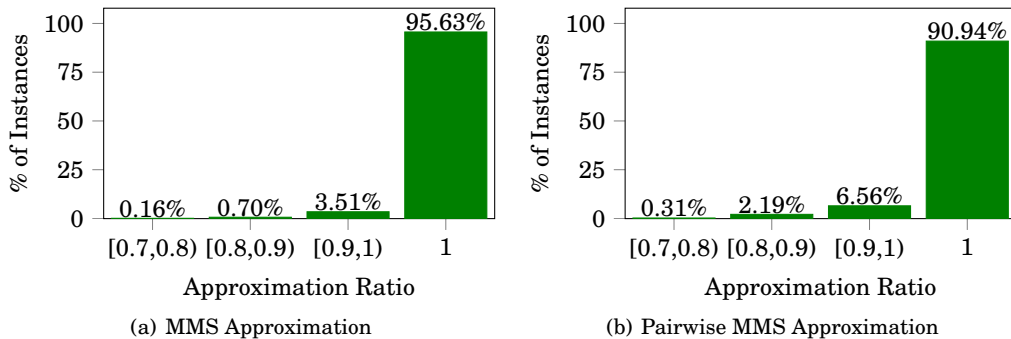


Fig. 1. MMS and Pairwise MMS approximation of the MNW solution on real-world data from Spliddit.

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 [Nguyen et al. 2013; Ramezani and Endriss 2010]. 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 [Nongillard et al. 2009] 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. In the full version, we show that S can be computed easily by finding a maximum cardinality matching in an appropriate bipartite graph. The problem then 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 1000 points among the goods. This in turn ensures that the utilities to the players will also be integral, and not more than 1000. 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 [Nguyen et al. 2013], 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:

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

Fig. 2. Nonlinear discrete optimization program

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

Fig. 4. MILP using segments on the log curve




Fig. 3. The log function and its approximations




Fig. 5. Running time of our MNW implementation

$\sum_{i \in \mathcal{N}} \sum_{t=2}^{1000} (\log t - \log(t-1)) \cdot U_{i,t}$, 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 .⁵ 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 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 1000. 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 2.9 GHz quad-core computer with 32

⁵In fact, this transformation is useful in maximizing any concave function, or minimizing any convex function, and thus may be of independent interest.

GB RAM, using CPLEX to solve the MILPs. The indicator-variables-based approach failed to run within our time limit (60 seconds) 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 30 seconds (whereas even the largest of the 1281 instances on Spliddit has 10 players). In fact, the algorithm solves every Spliddit instance in less than 3 seconds.

The largest real-world instance we have seen was actually reported offline by a Spliddit user. He needed to split an inheritance of roughly 1400 goods with his 9 siblings. Our implementation solves an instance of this size in roughly 15 seconds.

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 4.5).

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 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 paper is to advocate 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 (an example is presented in the full version). 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 (refer to the full version for a proof).

Throughout the paper 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. In the full version, we outline 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 the full version 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 strategyproof solutions. For example, Schummer [1997] shows that the only strategyproof 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).⁶ Second, we do not 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.

REFERENCES

- M. Aleksandrov, H. Aziz, S. Gaspers, and T. Walsh. 2015. Online Fair Division: Analysing a Food Bank Problem. In *Proc. of 24th IJCAI*. 2540–2546.
- G. Amanatidis, E. Markakis, A. Nikzad, and A. Saberi. 2015. Approximation Algorithms for Computing Maximin Share Allocations. In *Proc. of 42nd ICALP*. 39–51.
- K. J. Arrow and M.D. Intriligator (Eds.). 1982. *Handbook of Mathematical Economics*. North-Holland.
- H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. 2015. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence* 227 (2015), 71–92.

⁶In theory, one can hope to circumvent this result by making manipulation computationally hard [Bartholdi et al. 1989]. This is almost surely true (in the worst-case sense of hardness) for the MNW solution, where even computing the outcome is hard.

- J. Bartholdi, C. A. Tovey, and M. A. Trick. 1989. The Computational Difficulty of Manipulating an Election. *Social Choice and Welfare* 6 (1989), 227–241.
- A. Bogomolnaia and H. Moulin. 2001. A New Solution to the Random Assignment Problem. *Journal of Economic Theory* 100 (2001), 295–328.
- A. Bogomolnaia and H. Moulin. 2004. Random Matching under Dichotomous preferences. *Econometrica* 72 (2004), 257–279.
- S. Bouveret and J. Lang. 2008. Efficiency and envy-freeness in fair division of indivisible goods: logical representation and complexity. *Journal of Artificial Intelligence Research* 32 (2008), 525–564.
- S. Bouveret and M. Lemaître. 2016. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems* 30, 2 (2016), 259–290.
- S. J. Brams, M. Kilgour, and C. Klamler. 2015. Maximin Envy-Free Division of Indivisible Items. Manuscript. (2015).
- S. J. Brams and A. D. Taylor. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.
- E. Budish. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
- E. Budish, Y.-K. Che, F. Kojima, and P. Milgrom. 2013. Designing Random Allocation Mechanisms: Theory and Applications. *American Economic Review* 103, 2 (2013), 585–623.
- R. Cole and V. Gkatzelis. 2015. Approximating the Nash Social Welfare with Indivisible Items. In *Proc. of 47th STOC*. 371–380.
- E. Eisenberg and D. Gale. 1959. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics* 30, 1 (1959), 165–168.
- D. Foley. 1967. Resource allocation and the public sector. *Yale Economics Essays* 7 (1967), 45–98.
- A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. 2011. Dominant Resource Fairness: Fair Allocation of Multiple Resource Types. In *Proc. of 8th NSDI*. 24–37.
- J. Goldman and A. D. Procaccia. 2014. Spliddit: Unleashing Fair Division Algorithms. *SIGecom Exchanges* 13, 2 (2014), 41–46.
- F. P. Kelly. 1997. Charging and rate control for elastic traffic. *European Transactions on Telecommunications* 8 (1997), 33–37.
- D. Kurokawa, A. D. Procaccia, and N. Shah. 2015. Leximin Allocations in the Real World. In *Proc. of 16th EC*. 345–362.
- D. Kurokawa, A. D. Procaccia, and J. Wang. 2016. Can the maximin share guarantee be guaranteed?. In *Proc. of 30th AAAI*. Forthcoming.
- E. Lee. 2015. APX-Hardness of Maximizing Nash Social Welfare with Indivisible Items. arXiv:1507.01159. (2015).
- R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proc. of 6th EC*. 125–131.
- H. Moulin. 2003. *Fair Division and Collective Welfare*. MIT Press.
- J. Nash. 1950. The bargaining problem. *Econometrica* 18, 2 (1950), 155–162.
- T. T. Nguyen, M. Roos, and J. Rothe. 2013. A survey of approximability and inapproximability results for social welfare optimization in multiagent resource allocation. *Annals of Mathematics and Artificial Intelligence* 68, 1 (2013), 65–90.
- A. Nongillard, P. Mathieu, and B. Jaumard. 2009. A realistic approach to solve the Nash welfare. In *Proc. of 7th PAAMS*. 374–382.
- E. Pazner and D. Schmeidler. 1978. Egalitarian equivalent allocations: A new concept of economic equity. *Quarterly Journal of Economics* 92, 4 (1978), 671–687.
- A. D. Procaccia and J. Wang. 2014. Fair Enough: Guaranteeing Approximate Maximin Shares. In *Proc. of 14th EC*. 675–692.
- S. Ramezani and U. Endriss. 2010. Nash social welfare in multiagent resource allocation. In *Proc. of 12th AMEC*. 117–131.
- J. Schummer. 1997. Strategy-proofness versus efficiency on restricted domains of exchange economies. *Social Choice and Welfare* 14 (1997), 47–56.
- H. Varian. 1974. Equity, envy and efficiency. *Journal of Economic Theory* 9 (1974), 63–91.
- E. P. Wigner. 1960. The unreasonable effectiveness of mathematics in the natural sciences. *Communications on Pure and Applied Mathematics* 13, 1 (1960), 1–14.