Group Decision and Negotiation
https://doi.org/10.1007/s10726-024-09885-x

Two-Person Fair Division with Additive Valuations

D. Marc Kilgour[1] · Rudolf Vetschera[2]

Accepted: 13 March 2024

Abstract

In the literature, many desirable properties for allocations of indivisible goods have been proposed, including envy-freeness, Pareto optimality, and maximization of either the total welfare of all agents, the welfare of the worst-off agent, or the Nash product of agents' welfares. In the two-person context, we study relationships among these properties using both analytical models and simulation in a setting where individual preferences are given by additive cardinal utilities. We provide several new theorems linking these criteria and use simulation to study how their values are related to problem characteristics, assuming that utilities are assigned randomly. We draw some conclusions concerning the relation of problem characteristics to the availabilty of allocations with particular properties.

Keywords Fair division · Envy-freeness · Cardinal utilities

1 Introduction

How multiple independent participants can share a resource fairly is an important and easily understood social choice problem. Its importance is is illustrated by platforms such as Spliddit (Goldman and Procaccia 2015), which help users split an asset fairly. Is it possible to give each participant a satisfactory allocation, or is it impossible to reconcile their conflicting interests? (Thomson 2016) Many criteria for a good allocation have been proposed, including utilitarianism (Bentham 1789) and maximinality (Rawls 1971). Another criterion is based on envy; an agent envies another agent if the first prefers the second's share to its own. An allocation is envy-free if no agent envies any other (Gamow and Stern 1958; Kilgour and Vetschera 2018; Klamler 2021; Brams et al. 2023).

✉ Rudolf Vetschera
rudolf.vetschera@univie.ac.at

D. Marc Kilgour
mkilgour@wlu.ca

[1] Department of Mathematics, Wilfrid Laurier University, Waterloo, Canada

[2] Department of Business Decisions and Analytics, University of Vienna, Vienna, Austria

Published online: 18 April 2024

A resource that is continuously divisible can always be allocated in an envy-free way, although such an allocation may fail other desirable properties, such as contiguity (Procaccia 2016). But if the resource to be allocated consists of indivisible items, envy-freeness may simply be unachievable.

In this paper, we focus on the relationships between desirable properties of allocations such as, among others, envy-freeness and Pareto optimality. We perform this analysis in a specific setting: allocation of a finite set of indivisible, unsharable objects to two players whose preferences are known, measured on a cardinal scale and additive. This is clearly a very specific and restrictive model of preferences; it differs from the considerable part of the literature on fair division of indivisible items that assumes only that each player can rank the items. By studying what can be considered the extreme opposite setting, we attempt to delineate the possible relationships between allocation properties as far as possible. Relationships that cannot be established under these restrictive assumptions about preferences most likely will also not hold when considering more general models.

We study these relationships using a combined approach of formal proofs and large scale simulation. Results of this simulation not only confirm our analysis but also lead us to several conjectures which we leave for others to prove or disprove. They also provide insights into the likelihood that allocations exist with certain combinations of desirable properties. In other words, how often do various properties exist together? Furthermore, we are able to connect properties of allocations to properties of the problem such as the correlation between the players' utilities, offering additional insights into the mechanisms that make desirable properties feasible.

The remainder of the paper is structured as follows. Section 2 puts our work in the context of the related literature. In Sect. 3, we give an overview of our model of fair allocation of indivisible objects. Section 4 introduces the properties of allocations that we study, and Sect. 5 gives some analytical results on the relationships of these properties. An overview of the simulation model and its use is given in Sect. 6, and quantitative results deriving from the simulations are presented in Sect. 7. In Sect. 8, we summarize our results, discuss their limitations, and raise some open questions.

2 Related Literature

Whether an allocation of indivisible, unsharable items among several individuals is fair depends fundamentally on their preferences for subsets of items, or bundles. As already mentioned, the literature on fair division problems initially considered very general ordinal preference models in which players' rankings of individual items were expanded to rankings of bundles only via a cancellation condition—if the same (disjoint) bundle is added to each of two bundles, the ranking of the two bundles does not change. Consequently the rankings of bundles may be incomplete (e.g. Kilgour and Vetschera 2018; Bouveret et al. 2016; Brams et al. 2015, 2017). Nonetheless, Brams et al. (2003) used only these preference models in defining concepts such as envy-freeness and Pareto optimality, and obtained some results about their relationships.

More general preference models (including additive utilities, referred to as "one of the most classical settings in fair division of indivisible goods", p. 287) are comprehensively surveyed by Bouveret et al. (2016). However, the first reference to a model with additive utilities (Lipton et al. 2004) was dated 1 year after the survey of Brams et al. (2003). A similar survey that emphasizes connections to economics is Moulin (2019). Recent developments, with a focus on relaxations of fairness criteria and related algorithms, are covered in Amanatidis et al. (2023). The present paper follows the approach of Brams et al. (2023) but focuses on preferences that can be represented using additive utility functions.

The most often studied relationships in fair division problems are the trade-offs among various criteria of fairness and efficiency. As in many other branches of the study of collective decision making, this topic is a recurring theme. Caragiannis et al. (2012) formalized this trade-off using the *price of fairness*, a concept that appears frequently in the literature (e.g., Bouveret et al. 2016; Amanatidis et al. 2023). But attempts to use this concept to analyze the trade-off between envy-freeness and Pareto optimality run into a problem: Even in the case of only two players, it is possible that this trade-off cannot be calculated as no envy-free allocation exists (Bouveret et al. 2016).

The simplest instance is a "diamond and pebbles" allocation problem (Brams et al. 2023), where the set of objects to be distributed consists of one diamond and many pebbles, and for each player the utility of the diamond exceeds the total utility of all of the pebbles. Then no envy-free allocation is possible, as the diamond can be assigned to only one player. In fact, no envy-free allocation exists in any diamond-pebbles problem with an odd number of (approximately equally-valued) diamonds. As a consequence, the literature has considered relaxations of strict fairness criteria such as EF1, envy-freeness after removal of one item. The concept of price of fairness was also applied to such relaxations (Barman et al. 2020; Bei et al. 2021). Additional fairness criteria such as the Gini index and a gradual measure of envy have also been studied, e.g. by Aleksandrov et al. (2019).

In the present paper, we study the relationship between fairness and efficiency, using several concepts of efficiency including Pareto-optimality or (in the utilitarian sense) maximum utility sum. We also study the relationships among various fairness criteria in the two-player context, for example how often envy-free allocations also satisfy the Rawlsian criterion of maximizing the utility of the worse-off player.

Some solution concepts, such as the Nash bargaining solution (Nash 1950), aim to achieve both fairness and efficiency. Properties of the Nash bargaining solution are also widely studied in the fair division literature (Caragiannis et al. 2019). For additive utilities, Amanatidis et al. (2021) and Halpern et al. (2020) showed that the Nash solution is always EFX (envy-free after removal of any item) if valuations are binary (i.e., players evaluate items only as "good"or "bad"). This result was generalized to any additive preferences and EF1 by Suksompong (2023), who also showed that the Nash solution is the only additive solution with this property. Furthermore, the Nash solution approximates the Rawlsian Max–Min solution quite well (Caragiannis et al. 2019).

Our study uses simulation to analyze relationships among the properties of allocations. Although most works cited so far provide analytical results, there are also some simulation studies. Freeman et al. (2019) used synthetic data as well as actual preferences retrieved from the Spliddit platform to study the relationship between (approximations of) equality (each player receiving the same utility) and properties such as Pareto optimality. Dickerson et al. (2014) used a computational study to determine the threshold on the number of items at which EF allocations become likely to exist.

Our analysis extends the work of Brams et al. (2023) in several directions. We draw some general conclusions and then complement our analysis with a comprehensive simulation study. Our approach is to assume that both players' (additive) utilities are drawn at random, independently, from a distribution defined by Lebesgue measure.

3 The Model

We consider how a set of n indivisible items, $I = \{1, \dots, n\}$, can be shared by two players, A and B. An allocation $S = (S_A, S_B)$ is a partition of I into two subsets S_A and S_B, so that $S_A \cup S_B = I$ and $S_A \cap S_B = \emptyset$. In the allocation $S = (S_A, S_B)$, S_A is A's assignment or bundle, and S_B is B's. It is often convenient to describe an allocation by A's bundle, S_A; of course, $S_B = I \setminus S_A$, the complement of S_A.

We define $\mathcal{S}(I)$, or simply $\mathcal{S}$, to be the set of all possible allocations of the item set, I. By counting the possible assignments to A, it is clear that $|\mathcal{S}| = 2^n$.

We assume that player $X = A$ or B has utility $u_X(i) \geq 0$ for item $i \in I$, and that utilities are additive, so that player X's utility for the subset (or bundle) $S \subseteq I$ is

$$u_X(S) = \sum_{i \in S} u_X(i) \tag{1}$$

The utility of allocation $S = (S_A, S_B)$ to player X is $u_X(S) = u_X(S_X)$. We assume without loss of generality that utilities are scaled so that $u_A(I) = u_B(I) = 1$.

The utility of player X can be written as an n-vector

$$u_X = (u_X(1), u_X(2), \dots, u_X(n));$$

a pair of utility vectors (u_A, u_B) defines an allocation problem of size n. We assume that each of u_A and u_B is the realization of a uniform random process that produces non-negative numbers summing to 1, following a Lebesgue distribution. Moreover, we assume that u_A and u_B are independent. (Below we will describe how we generated such allocation problems in our simulation.)

The set of possible utility vectors for player X is

$$\left\{ (u_1, u_2, \dots, u_n) \, : \, u_i \geq 0 \; \forall \, i, \sum_{i=1}^{n} u_i = 1 \right\},$$

an $n-1$-dimensional set. A consequence of our assumptions is that any subset of utility vectors u_X with dimension strictly less than $n-1$ has probability (measure) zero; we ignore such subsets. For example, we assume that, if $1 \leq i < j \leq n$, then $0 < u_X(i) < 1$ and $u_X(i) \neq u_X(j)$ for $X = A$ or B. (Equalities like $u_X(i) = 0$ and $u_X(i) = u_X(j)$ can occur, but according to our probability model they can be neglected.) For the same reason, the probability is zero that a set $S \subseteq I$ exists such that $u_X(S) = u_X(I \setminus S)$. Similarly, if $S \neq \emptyset$, then $u_X(S) > 0$ and $u_X(I \backslash S) < 1$. The independence of u_A and u_B implies that, if S and T are both proper non-empty subsets of I, then $u_A(S) \neq u_B(T)$.

Given any allocation $S = (S_A, S_B)$, we take advantage of our principle that $u_A(S_A) \neq u_B(S_B)$ to call the player whose utility is less the weaker player, W, and write its utility as $u_W(S)$. Similarly, we call the player who receives more utility in allocation S the stronger player, G; its utility is $u_G(S)$.

Note that, for any allocation S, $W = W(S)$ and $G = G(S)$; in other words, the identity of the weaker and stronger player depends on the allocation. A consequence is that a statement like $u_W(S) < u_W(T)$ may compare the utilities of different players. For example, if $u(S) = (u_A(S), u_B(S)) = (0.2, 0.4)$ and $u(T) = (0.5, 0.3)$, then $u_W(S) = u_A(S)$ but $u_W(T) = u_B(T)$.

An allocation is envy-free (EF) for a player if its utility for the player's own bundle is at least as great as its utility for the opponent's bundle. An allocation is EF if it is EF for both players. Thus, an EF allocation $S = (S_A, S_B)$ satisfies $u_A(S_A) \geq u_A(S_B)$ and $u_B(S_B) \geq u_B(S_A)$. Using our assumptions that $u_X(S_A) + u_X(S_B) = 1$ for $X = A$ or B and that $u_X(S) \neq u_X(I \setminus S)$, a further simplification is possible: $S = (S_A, S_B)$ is EF iff $u_A(S_A) > 1/2$ and $u_B(S_B) > 1/2$.

These observations suggest a simple geometric view of allocations. Divide the unit square in (u_A, u_B)-space into four quadrants with lines at $u_A = 1/2$ and $u_B = 1/2$, and name the quadrants as follows:

I. $u_A > 1/2, u_B > 1/2$
II. $u_A > 1/2, u_B < 1/2$
III. $u_A < 1/2, u_B < 1/2$
IV. $u_A < 1/2, u_B > 1/2$

An allocation $S = (S_A, S_B)$ corresponds to the point $(u_A(S), u_B(S))$, which must lie within one of these quadrants. Allocation S is EF if and only if it corresponds to a point in Quadrant I.

Example 1 I contains $n = 4$ items and

$$u_A = (0.31, 0.22, 0.36, 0.11); \quad u_B = (0.23, 0.07, 0.38, 0.32)$$

The 16 possible allocations provide the players utilities given in Table 1. Figure 1 shows the available utilities for players A and B. Note that, as in any problem, the utility pairs (1, 0) and (0, 1) can be achieved by assigning all items to A or to B.

Observe that the $2^4 = 16$ allocations in Example 1 come in symmetric pairs. The allocation (S_A, S_B) corresponds to the point $(u_A(S_A), u_B(S_B))$, and the reverse allocation, (S_B, S_A) corresponds to $(u_A(S_B), u_B(S_A)) = (1 - u_A(S_A), 1 - u_B(S_B))$. In the Quadrant Diagram, these two points are symmetrically located with respect to (0.5, 0.5). (The line between them passes through (0.5, 0.5), and they are equally

Table 1 Utilities of all possible allocations in Example 1

(S_A, S_B)	$u_A(S_A)$	$u_B(S_B)$
({1,2,3,4}, ∅)	1	0
({1,2,3},{4})	0.89	0.32
({1,2,4},{3})	0.64	0.38
({1,3,4},{2})	0.78	0.07
({1,2},{3,4})	0.53	0.70
({1,3},{2,4})	0.67	0.39
({1,4},{2,3})	0.42	0.45
({1},{2,3,4})	0.31	0.77
({2,3,4},{1})	0.69	0.23
({2,3},{1,4})	0.58	0.55
({2,4},{1,3})	0.33	0.61
({2},{1,3,4})	0.22	0.93
({3,4},{1,2})	0.47	0.30
({3},{1,2,4})	0.36	0.62
({4},{1,2,3})	0.11	0.68
(∅,{1,2,3,4})	0	1

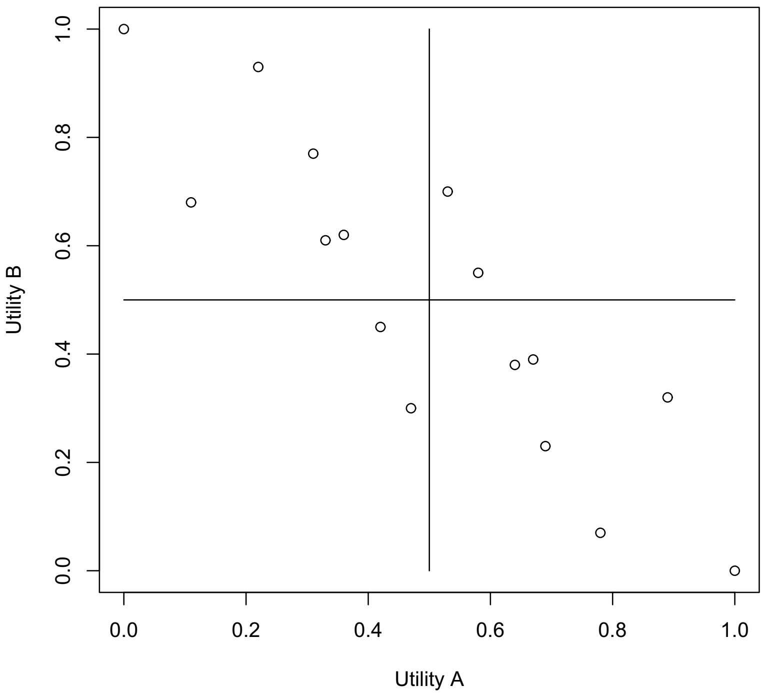

Fig. 1 Quadrant diagram of all possible allocations in Example 1

distant from this central point.) As a consequence, each allocation in Quadrant I is paired with an allocation in Quadrant III, and vice versa. Therefore, Quadrants I and III contain the same number of allocations. The relation of Quadrants II and IV is similar.

As already noted, any allocation in Quadrant I is EF, so Example 1 has two EF allocations, which are given by $(S_A, S_B) = (\{1,2\},\{3,4\}) = (12, 34)$ and $(S_A, S_B) = (23, 14)$. (When no confusion is possible, we simplify the notation for subsets in this way.) The reverse allocations, which lie in Quadrant III, are (34, 12) and (14, 23). Quadrants II and IV always contain some allocations—namely $(I, \emptyset)$ and $(\emptyset, I)$—and often contain most of them. For allocations in Quadrants II and IV, there is one-sided envy, which persists if the assignments are reversed—the only change is the identity of the envious player. For example, in a diamond and pebbles problem, all allocations lie near (1, 0), if *A* receives the diamond, or near (0, 1), if the diamond goes to *B*.

4 Properties of Allocations

4.1 Pareto and Lexicographic Optimality

It is a general principle of rationality that actors will never accept an outcome if another outcome that both prefer is available. Almost as compelling, at least to a social planner, is the principle that one outcome is better than another if both the weaker and stronger players do better—even if the identity of these two players is not the same.

Allocation *S* is Pareto-Superior (PS) to allocation *T* iff each player (weakly) prefers its assignment in *S* to its assignment in *T*, and at least one player strictly prefers its assignment in *S*. Because we ignore events that occur on sets of measure zero, for us allocation *S* is PS to allocation *T* iff each player strictly prefers its assignment in *S* to its assignment in *T*. In the Quadrant Diagram, *S* must lie northeast of *T*.

An allocation *T* is Pareto-Optimal (PO) iff there exists no allocation *S* such that *S* is PS to *T*. In the Quadrant Diagram, *T* is PO iff there is no allocation northeast of *T*. For Example 1, Fig. 1 shows that there are 8 PO allocations, only two of which are EF. Although it is not the case in Example 1, it is clearly possible for an EF allocation (in Quadrant I) not to be PO.

Any allocation *S* has an ascending utility vector, $(u_W(S), u_G(S))$, of which the entries are the players' utilities for their own assignments, written in ascending order. Allocation *S* is Lexicographically-Superior (LS) to allocation *T* iff each element of the ascending utility vector of *S* is at least equal to the corresponding element of the ascending utility vector for *T*, and in at least one case is strictly greater. Because we ignore the possibility of equality, for us *S* is LS to *T* if and only if $u_W(T) < u_W(S)$ and $u_G(T) < u_G(S)$. Allocation *T* is Lexicographically-Optimal (LO) if there exists no allocation *S* such that *S* is LS to *T*.

For Example 1, (13, 24) is PO (as Fig. 1 shows: $U_A(13) = 0.67$ and $U_B(24) = 0.39$). But is (13, 24) LO? The ascending utility vector of (13, 24) is (0.39, 0.67). But $u_A(12) = 0.53$ and $u_B(34) = 0.70$, so the ascending utility vector of

(12, 34) is (0.53, 0.70). Consequently, (12, 34) is LS to (13, 24), so (13, 24) is not LO.

For Example 1, we noted that (13, 24) is PO but not LO. As well (see Fig. 1), (1, 234) is PO but not LO—its utilities are (0.31, 0.77); its image, with utilities (0.77, 0.31) is Pareto-dominated by (123, 4) with utilities (0.89, 0.32). Note that each player's utility at the hypothetical allocation, S', is equal to the opponent's utility at S, so the probability associated with S' is zero—but, nonetheless, it links PO and LO.

4.2 Maximization

It stands to reason that the best allocation should maximize something (Klamler 2021; Brams et al. 2023). But what should it maximize? There have been several suggestions—the allocation should be Maximin (Rawls 1971), or maximize Total Welfare (Bentham 1789), or maximize Nash Welfare (Caragiannis et al. 2019; Nash 1950). Writing $(u_A(S_A), u_B(S_B)) = (x, y)$, we implement these ideas by searching for the allocation (S_A, S_B) that maximizes an appropriate function, $f(x, y)$.

An allocation is Maximin (MM) if there is no other allocation at which the minimum of the players' utilities is greater. Thus, a Maximin allocation selects (S_A, S_B) to maximize $f_{MM}(x, y) = \min\{x, y\} = \min\{u_A(S_A), u_B(S_B)\}$. We take the Maximin allocation to be unique, since two such allocations could occur only when two distinct subsets have equal utility sums. Thus the Maximin allocation is determined by

$$S_A^{MM} = \arg\max_{S_A \subseteq I} \left[\min\{u_A(S_A), u_B(I \setminus S_A)\}\right]$$

and $S_B^{MM} = I \setminus S_A^{MM}$. The *maximin utility* is

$$f_{MM}\left(u_A\left(S_A^{MM}\right), u_B\left(S_B^{MM}\right)\right) = \min\left\{u_A\left(S_A^{MM}\right), u_B\left(S_B^{MM}\right)\right\}.$$

In Example 1, $S^{MM} = (23, 14)$ achieves the Maximin utility, 0.55.

An allocation is Maximum Nash welfare (MNW) if there is no other allocation for which the product of the players' utilities is greater. Thus, a Maximum Nash welfare allocation is an allocation that maximizes the Nash Product $f_{MNW}(x, y) = x \cdot y$. Again, we take the MNW allocation to be unique, as another allocation with the same utility product could exist only on a set of measure zero. Thus the Maximum Nash welfare allocation is determined by

$$S_A^{MNW} = \arg\max_{S_A \subseteq I} \left[u_A(S_A) \cdot u_B(I \setminus S_A)\right]$$

and $S_B^{MNW} = I \backslash S_A^{MNW}$. In Example 1, $S^{MNW} = (12, 34)$, for which the Nash product equals 0.371.

An allocation is Maximum Total welfare (MTW) if there is no other allocation for which the sum of the players' utilities is greater. Thus, a Maximum Total welfare allocation maximizes $f_{MTW}(x, y) = x + y$. It is easy to calculate this allocation—just assign each item to the player who values it more. Again, we can take the MTW allocation to be unique, as another allocation with the same utility sum could exist only on a set of measure zero. Thus the Maximum Total welfare allocation is

$$S_A^{MTW} = \arg\max_{S_A \subseteq I} \left[u_A(S_A) + u_B(I \setminus S_A) \right]$$

and $S_B^{MNW} = I \backslash S_A^{MNW}$. In Example 1, $S^{MTW} = (12, 34)$, which has utility sum 1.23.

The level curves of the three functions f_{MM}, f_{MNW}, and f_{MTW} are compared in Fig. 2, where all curves pass through (0.6, 0.6). If this allocation is available, then each of these procedures would select it—providing there are no allocations above and to the right of the graph of the corresponding function. As the geometry makes clear, if $S^{MTW} = (0.6, 0.6)$, then $S^{MNW} = S^{MM} = (0.6, 0.6)$ also.

Example 1 illustrates that $S^{MM} \neq S^{MNW} = S^{MTW}$ can occur. As we will see below, all three allocations could be identical, or could differ from each other. The only restriction on the three values is given by Proposition 2 (see Sect. 5): if S^{MM} and S^{MTW} are equal, then they must both equal S^{MNW}. There is a sense in which S^{MM} and S^{MTW} are extremes; S^{MNW} is always "between" them.

Another point of comparison of S^{MM}, S^{MNW}, and S^{MTW} is their associated utility gaps. For any allocation S, the utility gap is $g(S) = u_G(S_G) - u_W(S_W)$, that is, the difference between the utilities (for their own assignments) of the stronger and the weaker player. (Recall that the player whose utility is greater is G, the stronger player; the player whose utility is less is W, the weaker player.) As we show in Proposition 3 (see Sect. 5), the utility gaps reflect the ordering of the three maximization outcomes.

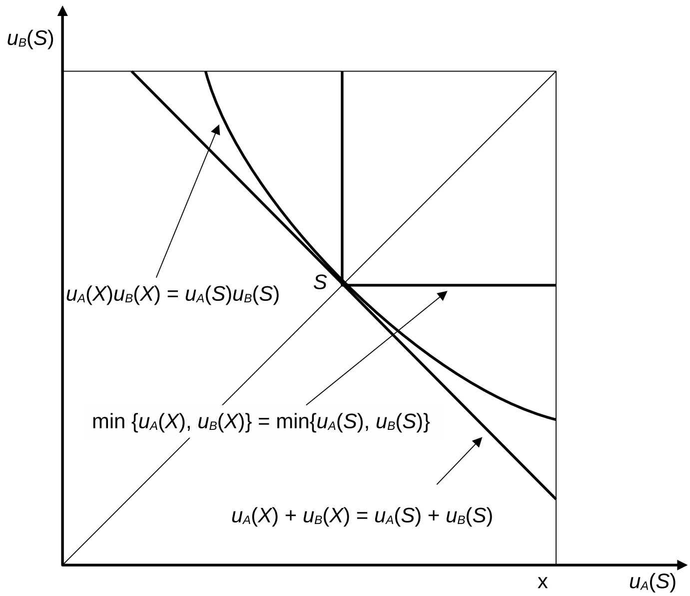

Fig. 2 Three functions to be maximized

4.3 Envy-freeness

As defined in Sect. 3, an allocation is envy-free for a player if and only if that player's utility for its own bundle is at least as great as its utility for the opponent's bundle. Moreover, an allocation is envy-free (EF) if and only if it is envy-free for both players. We noted that an allocation $S = (S_A, S_B)$ is EF if and only if $u_A(S_A) > 1/2$ and $u_B(S_B) > 1/2$.

This picture, that an EF allocation exists if and only if there is a maximin allocation S^{MM} satisfying $\min\left\{S_A^{MM}, S_B^{MM}\right\} > 1/2$ is the content of a Theorem of Brams et al. (2023), who showed that an EF allocation exists if and only if some maximin allocation is EF. In view of our assumptions, we conclude that an EF allocation exists if and only if $\min\left\{u_A\left(S^{MM}\right), u_B\left(S^{MM}\right)\right\} > 1/2$, which is to say that S^{MM} is EF, and thus lies in Quadrant I. Therefore, if S^{MM} is not EF, then there are no EF allocations, Quadrant I is empty (and so therefore is Quadrant III); every allocations fall into Quadrant II or Quadrant IV, where there is one-sided envy.

A weakening of envy-freeness is also of interest. If an allocation is not EF because, say, A envies B, then it is envy-free after the removal of any item (EFX) if A would not envy B were any single item removed from B's assignment. (Of course, the roles of A and B may be interchanged.) Note that the condition requires that A not envy B if any item assigned to B is removed from B'a bundle. Also, the item removed simply disappears from the calculation—it is not reassigned to A.

Brams et al. (2023) prove that if S^{MM} is not EF, then it is EFX. Below—Proposition 4 in Sect. 5—we present a geometric proof of this result. Observe that if there are no EF allocations S^{MM} must lie in Quadrant II or IV.

We have already noted that, if an EF allocation exists, then S^{MM} must be EF. In particular, if S^{MNW} or S^{MTW} is EF, then S^{MM} must be EF. Corollary 3 in Sect. 5 establishes the related fact that, if S^{MTW} is EF, then S^{MNW} is also.

5 Propositions

The two properties of Pareto Optimality and Lexicographic Optimality are closely related. LO is a stronger condition than PO—in a sense, it is twice as strong—as Proposition 1 makes clear. The relationship between PO and LO is illustrated in Fig. 3. In utility space, any allocation that is PS to S must lie in the shaded region above and to the right of S, so S is PO exactly when there are no allocations in this region. Allocations that are LS to S are located either in the PS region or the lightly shaded region that is the PS region for S', the point in utility space that is the reflection of S in the 45-degree line. Allocation S is LO if and only if no allocation is PS to S or to S'. Note that the two PO regions overlap in the upper right corner.

Proposition 1 *An allocation $T = (T_A, T_B)$ is LO if and only if T is PO and the allocation T', with utilities $\left(u_A\left(T'_A\right), u_B\left(T'_B\right)\right) = \left(u_B(T_B), u_A(T_A)\right)$, is also PO.*

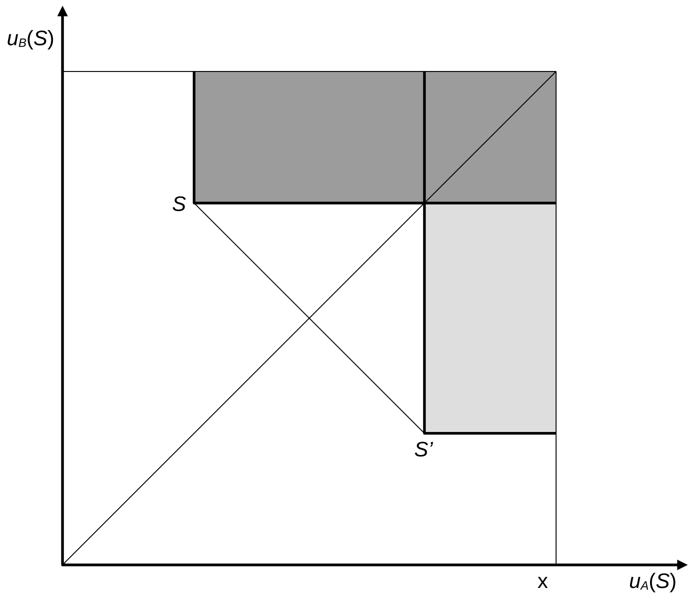

Fig. 3 Illustration of the relationship between LO and PO

Proof It is clear that, if T is LO, then T is PO, and so is T'. To show the converse, suppose that T is not LO, so that there exists an allocation $S = (S_A, S_B)$ that is LS to T. We show that either T is not PO or T' is not PO.

If the ascending utility vector of T is (x_1, x_2) and the ascending utility vector of S is (y_1, y_2), then $x_1 < y_1$ and $x_2 < y_2$. Suppose without loss of generality that $u_A(T_A) < u_B(T_B)$, so that the ascending utility vector of T is $(x_1, x_2) = \left(u_A(T_A), u_B(T_B)\right)$. If $u_A(S_A) < u_B(S_B)$, then the ascending utility vector of S is $(y_1, y_2) = \left(u_A(S_A), u_B(S_B)\right)$, so $u_A(T_A) < u_A(S_A)$ and $u_B(T_B) < u_B(S_B)$, proving that S is PS to T, and therefore that T is not PO.

Alternatively, it must be the case that $u_B(S_B) < u_A(S_A)$, so that the ascending utility vector of S is $(y_1, y_2) = (u_B(S_B), u_A(S_A))$. Consider an allocation T' with utilities as specified. Because $u_B(T'_B) = u_A(T_A) < u_B(T_B) = u_A(T'_A)$, the ascending utility vector of T' is $\left(x'_1, x'_2\right) = \left(u_B(T'_B), u_A(T'_A)\right)$. Now the fact that S is LS to T implies that $u_B(T'_B) < u_B(S_B)$ and $u_A(T'_A) < u_A(S_A)$. Therefore, S is PS to T', and T' is not PO, as required. □

We consider the allocation T' to be hypothetical—by assumption, it exists with probability zero because A's utility for T'_A equals B's utility for T_B.

As the next corollary shows, LO is a common property of the three maximization outcomes.

Corollary 1 *S^{MM}, S^{MNW}, and S^{MTW} are all LO.*

Proof As Fig. 4 shows, if S^{MM} is not LO, then there is an allocation in the shaded region, by Proposition 1. But this allocation, (S_A, S_B), must satisfy $\min\left\{u_A(S_A), u_B(S_B)\right\} > \min\left\{u_A\left(S_A^{MM}\right), u_B\left(S_B^{MM}\right)\right\}$, which is impossible by the definition of S^{MM}. Therefore there is no allocation in the shaded region, and S^{MM} is LO. The same argument shows that S^{MNW} and S^{MTW} are LO. □

We now consider whether the three maximization outocmes can be identical or different. Our first observation is the next Proposition.

Proposition 2 *If* $S^{MM} = S^{MTW}$, *then* $S^{MM} = S^{MNW} = S^{MTW}$.

Proof Let $\left(S_A^0, S_B^0\right)$ be the allocation S^{MM}, and let $x_0 = u_A\left(S_A^0\right)$ and $y_0 = u_B\left(S_B^0\right)$. Assume without loss of generality that $x_0 < y_0$. Suppose that there exists $S_A^1 \subseteq I$ such that, if $x_1 = u_A\left(S_A^1\right)$ and $y_1 = u_B\left(I \backslash S_A^1\right)$, then $x_1 y_1 > x_0 y_0$ but $\min\{x_1, y_1\} < \min\{x_0, y_0\}$. Thus $S^{MNW} \neq S^{MM}$. We show that it cannot be the case that $S^{MTW} = S^{MM}$.

Assume that $x_1 < x_0$. Then, because $x_1 y_1 > x_0 y_0$, $y_1 > y_0$. For $y \in [y_0, 1]$, define $h(y) = \frac{x_0 y_0}{y}$ and $k(y) = x_0 + y_0 - y$. Then $h(y_0) - k(y_0) = 0$ and, by calculus, $h(y) - k(y)$ is increasing in y, so $h(y_1) > k(y_1)$. Now $x_1 y_1 > x_0 y_0$ implies that $x_1 > h(y_1)$, so it follows that $x_1 > k(y_1)$, and therefore that $x_1 + y_1 > x_0 + y_0$. In

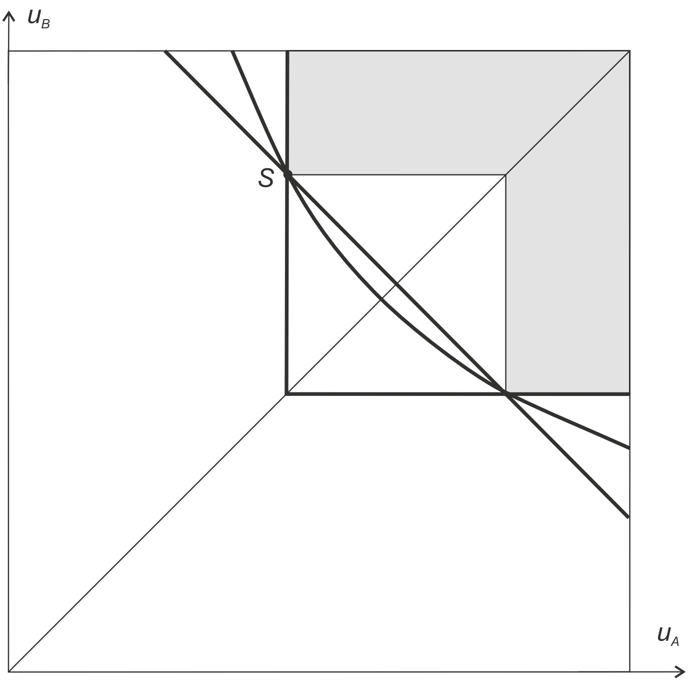

Fig. 4 Proof of Corollary 1

particular, (x_0, y_0) does not maximize f_{MTW}, so that $S^{MTW} \neq S^{MM}$. The proof is similar if $x_0 < x_1$. □

Figure 5 illustrates Proposition 2. The hypothesis that $S^{MM} = S^{MTW}$ implies that there are no allocations in the shaded region. It follows that there are no allocations on any level curve of $f_{MNW}(x, y)$ that lies above and to the right of $S^{MM} = S^{MTW}$, so $S^{MNW} = S^{MM} = S^{MTW}$. In consequence, if S^{MM}, S^{MNW}, and S^{MTW} are not identical, then S^{MM} and S^{MTW} must be different.

Recall that the utility gap of an allocation $S = (S_A, S_B)$ is $g(S) = |u_A(S_A) - u_B(S_B)| = u_G(S_G) - u_W(S_W)$. The utility gaps reinforce the suggestion of Proposition 2 concerning the ordering of maximization points.

Proposition 3 *The utility gaps of S^{MM}, S^{MNW}, and S^{MTW} satisfy*

$$g\left(S^{MM}\right) \leq g\left(S^{MNW}\right) \leq g\left(S^{MTW}\right).$$

Moreover, equality holds only if the corresponding allocations are identical.

Proof Clearly, the utility gaps of identical allocations are equal, so we need only show that the inequalities are strict when the corresponding allocations are different. Suppose that $S^{MM} \neq S^{MNW}$. Let $\left(u_W\left(S^{MM}\right), u_G\left(S^{MM}\right)\right) = (x_0, y_0)$, $\left(u_W\left(S^{MNW}\right), u_G\left(S^{MNW}\right)\right) = (x_1, y_1)$, and $(x_0, y_0) \neq (x_1, y_1)$. Then, assuming without

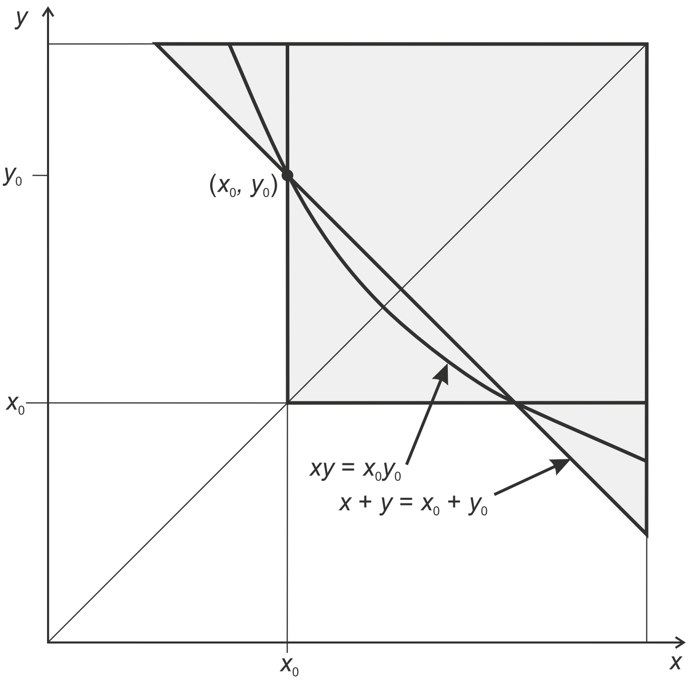

Fig. 5 Proof of Proposition 2

loss of generality that $x_0 < y_0$ and $x_1 < y_1$, it follows that $x_1 y_1 > x_0 y_0$ and $x_1 < x_0$. Therefore

$$g(S^{MNW}) = y_1 - x_1 > \frac{x_0 y_0}{x_1} - x_1 > \frac{x_0 y_0}{x_1} - x_0 > y_0 - x_0 = g(S^{MM})$$

Now assume that $S^{MNW} \neq S^{MTW}$, and let $\left(u_W\left(S^{MTW}\right), u_G\left(S^{MTW}\right)\right) = (x_2, y_2)$. Again using $\left(u_W\left(S^{MNW}\right), u_G\left(S^{MNW}\right)\right) = (x_1, y_1)$, define $d_x = x_1 - x_2$ and $d_y = y_1 - y_2$. As usual, we can assume without loss of generality that $x_2 < y_2$. By the definition of S^{MTW}, $d_x + d_y = (x_1 + y_1) - (x_2 + y_2) < 0$, so at least one of d_x and d_y is negative.

From the definition of S^{MNW}, it follows that

$$(x_2 + d_x)(y_2 + d_y) > x_2 y_2 \tag{2}$$

Condition (2) must fail if both d_x and d_y are negative, so we conclude that exactly one of d_x and d_y is negative.

To complete the proof, we show that $d_x > 0$ and $d_y < 0$. Assume, to obtain a contradiction, that $d_y > 0$. From (2), it follows that

$$x_2 d_y + y_2 d_x + d_x d_y > 0 \tag{3}$$

But the assumption that $d_y > 0$ implies that $x_2 d_y + y_2 d_x < y_2 d_y + y_2 d_x = y_2(d_x + d_y) < 0$. Because $d_x + d_y < 0$, this conclusion contradicts (3), so we conclude that $d_y < 0$ must hold. We conclude that

$$g\left(S^{MNW}\right) = y_1 - x_1 = y_2 + d_y - (x_2 + d_x) = y_2 - x_2 + (d_y - d_x) \leq g\left(S^{MTW}\right) \tag{4}$$

since $g\left(S^{MTW}\right) = y_2 - x_2$. □

A graphical proof of the first part of Proposition 3 is shown in Fig. 6. If there is a point on a level curve of f_{MNW} that lies above and to the right of the level curve of f_{MNW} that contains (x_0, y_0), and that is not MM (i.e., has a lower u_W than S^{MM}), this point must lie in one of the two shaded regions. In both shaded regions, the utility gap is greater than $y_0 - x_0$.

Figure 7 gives a graphical proof for the second part of Proposition 3: if $S^{MNW} \neq S^{MTW}$, then $g\left(S^{MNW}\right) < g\left(S^{MTW}\right)$. Any point (x_2, y_2) for which $x_2 \cdot y_2 < x_1 \cdot y_1$ and $x_2 + y_2 > x_1 + x_1$ must lie in one of the shaded regions, where $g(x_2, y_2) > g(x_1, y_1)$.

From the proof of Proposition 3, we immediately obtain

Corollary 2 $u_W\left(S^{MM}\right) \geq u_W\left(S^{MNW}\right) \geq u_W\left(S^{MTW}\right)$. *Moreover, equality holds only if the corresponding allocations are identical.*

An allocation $S = (S_A, S_B)$ is EF if and only if $u_W(S_W) > 1/2$, which implies the next result.

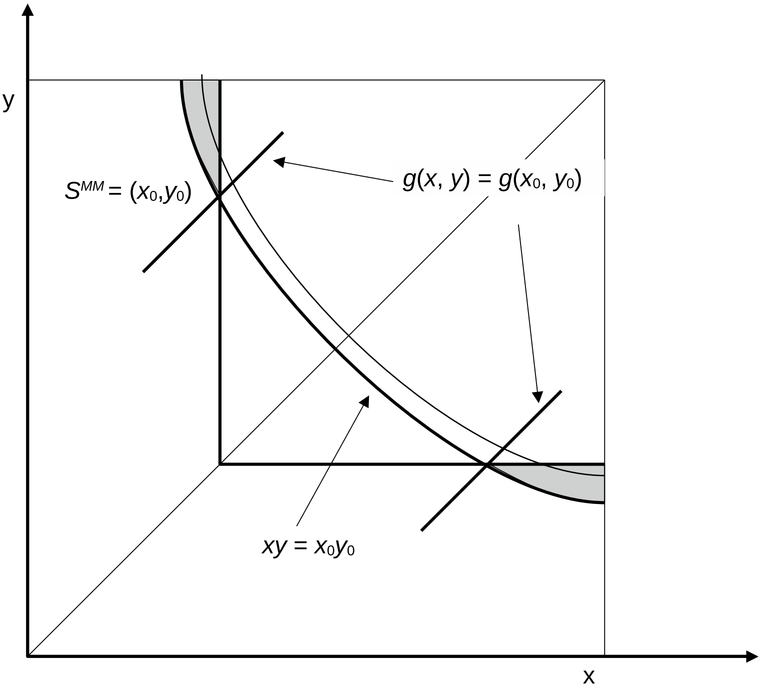

Fig. 6 Proof of Proposition 3a

Corollary 3 *If S^{MNW} is EF, then S^{MM} is EF. If S^{MTW} is EF, then both S^{MNW} and S^{MM} are EF.*

We are able to give a new proof to the Theorem of Brams et al. that applies when there is no EF allocation.

Proposition 4 (Brams et al. 2023): *If S^{MM} is not EF, then it is EFX.*

Proof First observe that if there is no EF allocation, then every allocation must lie in Quadrant II or Quadrant IV. If S^{MM} is not EF and if $x = u_A\left(S_A^{MM}\right)$, then we can assume without loss of generality that $x < u_A\left(S_B^{MM}\right) = 1 - x$ and therefore that $x < 1/2 < 1 - x$.

Now define $s = \min_{t \in S_B^{MM}} \{u_A(t)\}$. We will show that

$$x > 1 - x - u_A(s), \tag{5}$$

which implies that A does not envy B after s is removed, and therefore that S^{MM} is EFX.

We will prove (5) by contradiction. If it fails, we must have $x + u_A(s) < 1 - x$. Define the allocation S' by $S'_A = S_A^{MM} \cup \{s\}$ and

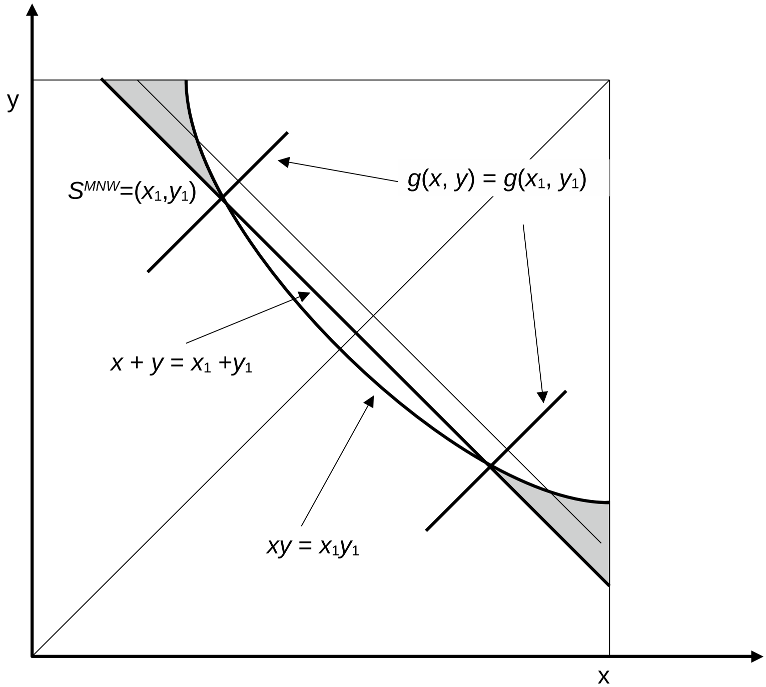

Fig. 7 Proof of Proposition 3b

$S'_B = S_B^{MM} \setminus \{s\}$. Suppose first that $u_B(S'_B) = u_B(S^{MM}) - u_B(s) > \frac{1}{2}$. Then $u_A(S_A^{MM}) < x + u_A(s) = u_A(S'_A) < \frac{1}{2} < u_B(S'_B)$, contradicting the fact that S^{MM} is a maximin allocation. Therefore, if (5) fails, we must have $u_B(S'_B) = u_B(S^{MM}) - u_B(s) < \frac{1}{2}$. But this shows that allocation S' lies in Quadrant III, which is impossible, as already noted. We conclude that (5) is true, which implies that S^{MM} is EFX. □

To summarize, the three maximization outcomes, S^{MM}, S^{MNW}, and S^{MTW}, which can be taken to be unique, are always LO. If they are not all the same, then S^{MM} and S^{MTW} must be different, in which case S^{MNW} may equal S^{MM}, or may equal S^{MTW}, or may be different from both of them. If S^{MM} is not EF, then neither are the other two. If S^{MNW} is EF, then so is S^{MM}. If S^{MTW} is EF, then so are the other two. We now use simulation to study how often these events occur.

6 Simulation Model

To study the relationships and frequencies of the properties of allocations introduced in Sect. 4, we conducted a comprehensive set of computer simulations of two-person allocation problems of various sizes. Each problem instance was defined by two randomly generated utility vectors, one for each party. We used the method of Butler et al. (1997) to generate vectors according to a uniform distribution over the unit simplex of dimension $n-1$. For each instance, we generated all possible allocations of items and then determined their properties. Due to the limited precision of digital computers, items or bundles occasionally had equal utilities within numerical tolerance. We discarded all problems in which the MM, MNW or MTW allocations were not unique, which happened in less than 0.00001% of all problems generated.

The simulation program was written in Pascal using the open source Free Pascal compiler version 3.2.2 and the Lazarus development environment (available from https://www.lazarus-ide.org/). The program code is available upon request from the authors.

We performed simulations for problems with $n = 4, 5, \dots,$ and 12 items, carrying out two sets of experiments.

(1) We generated 50,000,000 problems of each size, considered all possible allocations, and recorded summary data on the numbers exhibiting each of the properties EF, EFX, MM, MNW, MTW, PO, and LO, as well as combinations of these properties. In particular, we counted the numbers of EF and EFX allocations, and also recorded whether a problem allowed for any such allocation. Formally, four combinations are possible but, as Proposition 4 (Brams et al. 2023) shows, a problem with no EF allocation must have an EFX allocation, so only three combinations can actually occur. We also classified problems according to the MM, MNW and MTW allocations. (As already indicated, we discarded problems in which these allocations were not unique.) Our classification was based on whether these three outcomes were the same or different. Because of Proposition 2, a problem must fall into exactly one of the four classes shown below:

M1N1T1 One allocation has all three properties: $S^{MM} = S^{MNW} = S^{MTW}$.
M1N1T2 One allocation is MM and MNW, but the MTW allocation is different: $S^{MM} = S^{MNW} \neq S^{MTW}$.
M1N2T2 One allocation is MM, and a different allocation is both MNW and MTW: $S^{MM} \neq S^{MNW} = S^{MTW}$.
M1N2T3 All three allocations are different: $S^{MM} \neq S^{MNW} \neq S^{MTW}$ and $S^{MM} \neq S^{MTW}$.

To understand our notation, note that the three properties correspond to the letters M, N and T. The number following each letter indicates whether the property occurs at the same allocation as the previous property, or a different allocation.

(2) We generated 1,000,000 problems of each size, recording the utility vectors, the S^{MM}, S^{MNW}, and S^{MTW} allocations, as well as other properties. Thus, we could

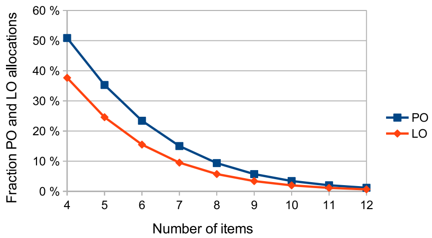

Fig. 8 Fractions of allocations that are PO and LO

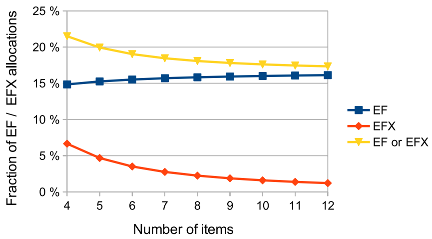

Fig. 9 Fractions of allocations that are EF and EFX, and problems without EF allocations

study the impact of problem characteristics, such as the correlation of utilities, on properties of allocations. Data storage requirements implied a smaller sample size, but one that is large enough that statistical conclusions could be fairly precise.

7 Quantitative Results

7.1 Efficiency and Fairness Properties

We first consider the fraction of allocations exhibiting individual properties. Since the S^{MM}, S^{MNW} and S^{MTW} outcomes are unique, we focus on the properties PO, LO, EF, and EFX.

Figure 8 shows two views of the fraction of PO and LO allocations as a function of number of items. The decrease in these fractions is greater than exponential, as evidenced by the logarithmic scale of the right-hand graph.

The left-hand graph of Fig. 9 shows the fraction of all allocations that are EF or EFX. The right-hand graph of Fig. 9 shows the fraction of problems with no EF allocation. This share drops exponentially; more than 99% of problems with 6 items have at least one EF allocation. Below we conclude that, in practice, lack of EF allocations is a concern only for small problems.

Perhaps the most interesting observation in Fig. 9 is the disappearance of EFX allocations as the number of items increases. Two factors must come together in order for an allocation that is not EF to be EFX: The utility difference between the two players' bundles must be small enough that removing just one item from the envied player reverses the envious player's utilities, and the utility of the envied player's least preferred (to the envious player) item must be large enough to cover this gap. In expectation, the utility of a least preferred item decreases in the number of items, and the expected difference in utilities also decreases. Our simulation indicates that the former effect is stronger, so that EFX allocations become less common as the number of items increases.

In contrast, the fraction of EF allocations increases in the number of items, but this increase does not compensate for the decrease in EFX allocations, so the total fraction of allocations that are either EF or EFX decreases. We have no theoretical explanation for the observed tendency of the total number of EF allocations to approach a limit as the number of items increases, so we formulate it as a conjecture.

Conjecture 1 *The fraction of all allocations that are EF converges to about* 1/6 *as* n *increases*.

Next, we consider the interaction between properties PO and EF. Neither of these proprieties implies the other, so it is possible that allocations exists which exhibit none, only one, or both properties. Figure 10 shows how the share of allocations among these four groups develops as the number of items increases.

The fraction of allocations that are neither PO nor EF first increases sharply, and then levels off at around 5/6 of all allocations. Only a relatively small fraction of allocations (about 0.5% in problems with 12 items) are both PO and EF, illustrating the difficulty of identifying allocations with desirable properties. For smaller problems, the number of allocations that are PO but not EF exceeds the number of allocations that are EF but not PO; this relation is reversed for larger problems. Essentially the same observations apply to the study of LO versus EF, so they are not reported here.

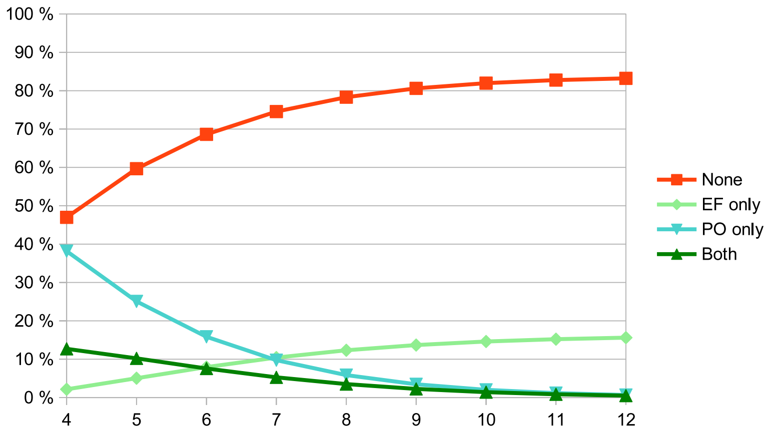

Fig. 10 Fraction of allocations that are EF and/or PO

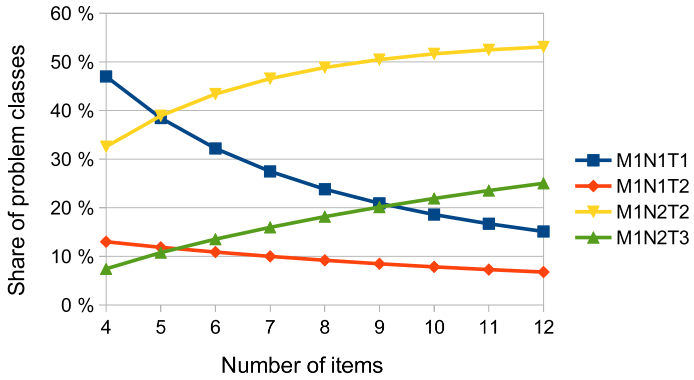

Fig. 11 Frequency of problem classes for different problem sizes

7.2 Maximization properties

Because it had unique MM, MNW,and MTW allocations, each problem fell into one of the four classes defined above. Figure 11 shows the frequency of these classes as a function of problem size. Perhaps it is not surprising that, as the number of items increases, problems in which a single allocation has all three properties (class M1N1T1) become less frequent and problems in which all three outcomes are different (class M1N2T3) become more frequent. The other two classes develop differently. Class M1N2T2, in which MNW and MTW occur at the same allocation while the MM allocation is different, becomes the largest class as the problem size passes five items. Meanwhile, M1N1T2, the class in which MM and MNW are the same but MTW is different, becomes the smallest class at about the same point. We connect this observation to the view of Rachmilevitch (2015) that the Nash bargaining solution tends to emphasize efficiency over fairness, so it is more commonly close to MTW than to MM.

Next, we analyze how frequently the three allocations S^{MM}, S^{MNW}, and S^{MTW} are envy-free, performing this analysis separately for each problem class. As Fig. 12 makes clear, the fraction of problems in which all maximizing allocations are EF increases with problem size, a statement that is true of every problem class. In class M1N1T1, where one allocation has all three properties, this allocation must be EFX if it is not EF, according to Proposition 4. In the other three problem classes, allocations that are different from S^{MM} might be EF or EFX or neither of the two. (For clarity, EFX is shown only for class M1N1T1; for the other classes, we only distinguish between EF and non-EF allocations.)

In fact, in the two classes in which $S^{MNW} = S^{MTW}$, this allocation is almost always EF. As well, when the size is larger, problems that lack an EF allocation become rare: our simulation data found none of them of size at least 11 items (although we believe that such problems exist). But when $S^{MNW} \neq S^{MTW}$, it was fairly common that neither of these allocations is EF, particularly for larger problem sizes. In summary, the probability that allocations focusing on efficiency are not EF increases as the number of items increases.

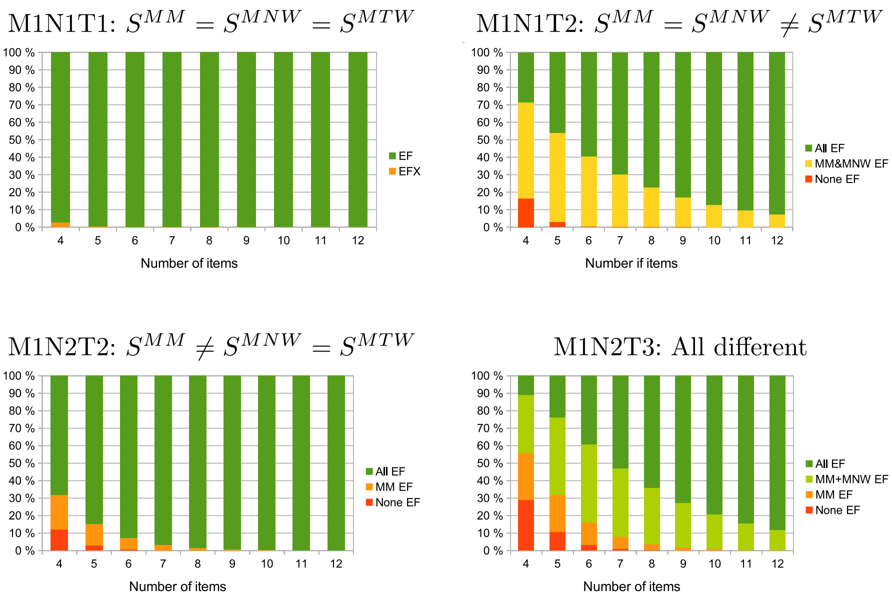

Fig. 12 Fractions of cases in which Maximization outcomes are EF

7.3 Problem characteristics

Finally, we analyze the relationship between problem characteristics and the properties of allocations. We focus on the correlation of the parties' utilities, which is a natural measure of the level of conflict: In the limit, when the correlation of the players' utilities is -1, an allocation that assigns every item to the player who prefers it is MM, MNW and MTW. If utilities are positively correlated, both players have high values for the same items, so the allocation problem is more difficult. We performed similar analyses using other measures such as the minimum utility of any item, or the maximum absolute difference in utilities over all items, or the number of items preferred by each player. These measures led to qualitatively similar results, so for brevity we present only the results for correlation.

Figure 13 shows the relationship between the number of EF allocations in a problem and the utility correlation. In this and the following figures, we present only the results for our smallest ($n = 4$) and largest ($n = 12$) problems. (Results for all other problem sizes are available upon request. In general, they follow the trend between these two extreme points.)

It is clear that there is a negative relation between utility correlation and the number of EF allocations, at least up to a certain point. All 4-item problems with no EF allocation are positively correlated. Similarly, 12-item problems with fewer than 100 EF allocations (the minimum in the sample was 2) occur only when the utility correlation exceeds about 0.7. However, this relationship is not deterministic. For example, we found problems with 4 items and correlation exceeding 0.9 that nonetheless exhibited 4 EF allocations, and 12-item problems with correlation exceeding

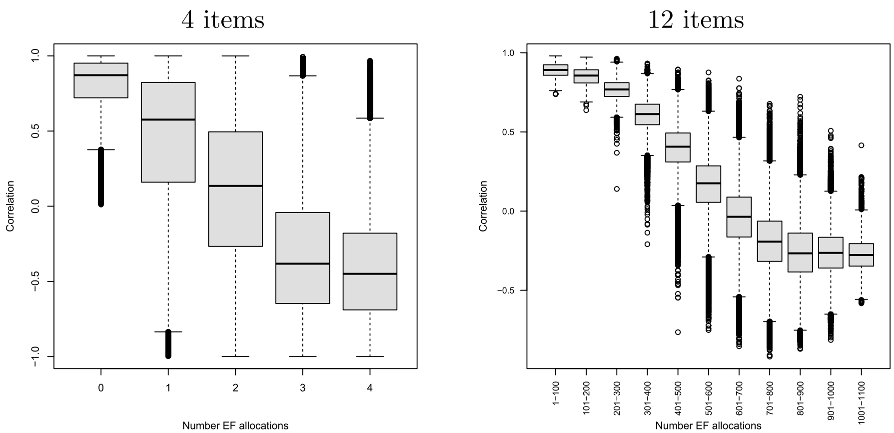

Fig. 13 Distribution of utility correlations in problems with different numbers of EF allocations

0.7 that had more than 800. At about 3/4 of the maximum number of EF allocations, this relationship flattens: the dispersion of correlation coefficients decreases, but the average and median values remain almost the same.

For every value of *n*, we observed a maximum of 2^{n-2} EF allocations. There is no doubt that a problem can have 2^{n-2} EF allocations. For example, suppose that each player has an item with utility more than 1/2, and that these two items are different. Then an allocation is EF if and only if it assigns to each player its most valued item. There are 2^{n-2} such allocations. However, our simulation data contained problems with exactly 2^{n-2} EF allocations that did not contain two such high-valued items, so we consider that we lack a full understanding of this apparent upper limit. We formulate another conjecture:

Conjecture 2 *The maximum number of envy-free allocations in a problem with n items is* 2^{n-2}.

Note that this conjecture depends on our assumptions of additive cardinal preferences and unique utility values for all items and bundles. Without these assumptions the conjecture can fail. For example, consider a 4-item problem in which both players are indifferent among all items (i.e., both players have utility 1/4 for each item). Clearly, every allocation that assigns two items to each player is EF; there are $\binom{4}{2} = 6 > 4$ such allocations.

Next we consider the relationship between the problem classes and utility correlation. As Fig. 14 shows, when utility correlation is negative, most problems with four items and many with 12 fall into class M1N1T1 (all three maximization principles yield the same outcome). These are the same problems that tend to have a large number of EF allocations. In these cases, the unique maximization outcome is not necessarily EF, though it must be PO (in fact, LO), and the number of PO (and LO) allocations increases with the correlation. At the other extreme, when the players' evaluations are identical, class M1N2T3 is most likely; moreover, all allocations are

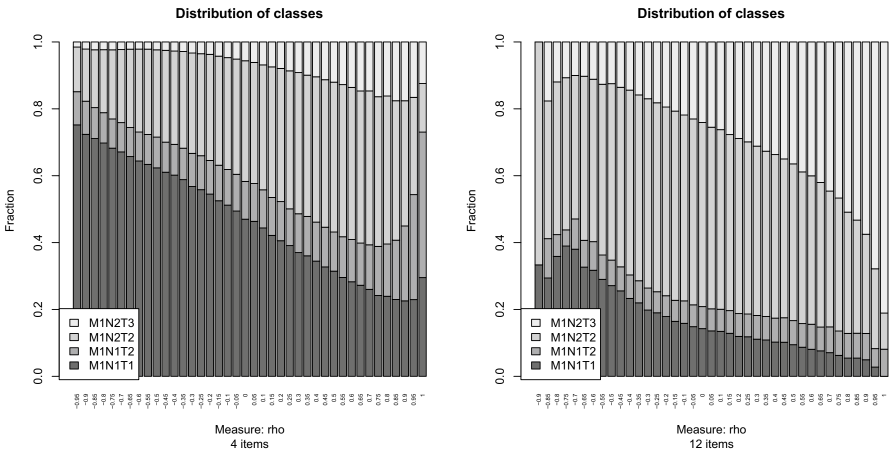

Fig. 14 Distribution of problem classes for different levels of utility correlation

PO. Thus in problems with a positive utility correlation, especially if it is high, the three maximization outcomes occur at different allocations. Again, we note that the relationship is not deterministic; in all problem classes, there are instances of utility correlation close to +1 and instances where it is close to −1.

Finally, Figs. 15 ($n = 4$) and 16 ($n = 12$) show the impact of correlation on whether the allocations S^{MM}, S^{MNW}, and S^{MTW} are EF. Again, the relationship is clear. If $n = 4$, all problems in which no maximization allocation is EF exhibit a positive utility correlation. If $n = 12$, this observation also holds for problems in which only S^{MM} is EF. (Our sample contained no problems of that size without EF allocations.) Problems in which the other allocations are also EF on average have a lower correlation of utilities, although this phenomenon occurs in occasional problems with correlation close to 1.

To further study the relationship between utility correlation and the envy-freeness of maximizing allocations, we calculated the minimum correlation coefficients of problems in which the three maximizing allocations are EF or not. The results of this analysis are presented in Table 2. The most striking result in this table is that except for the case of five items, there are no problems with negatively correlated utilities in which S^{MM} is not envy-free. But we know that if any allocation is EF, then S^{MM} must be EF, so this observation implies that if the number is items is different from five, then no EF allocation exists only when utilities are positively correlated.

To confirm this observation, we performed a simulation with three items and found several instances in which a problem with negative utility correlation had no EF allocation. We then performed additional simulations for 4-item and 7-item problems, 15 million problems each. In all of these simulations, an EF allocation was found whenever utility correlations were negative. The smallest correlation coefficients in problems without EF allocations were 0.320 for problems with seven

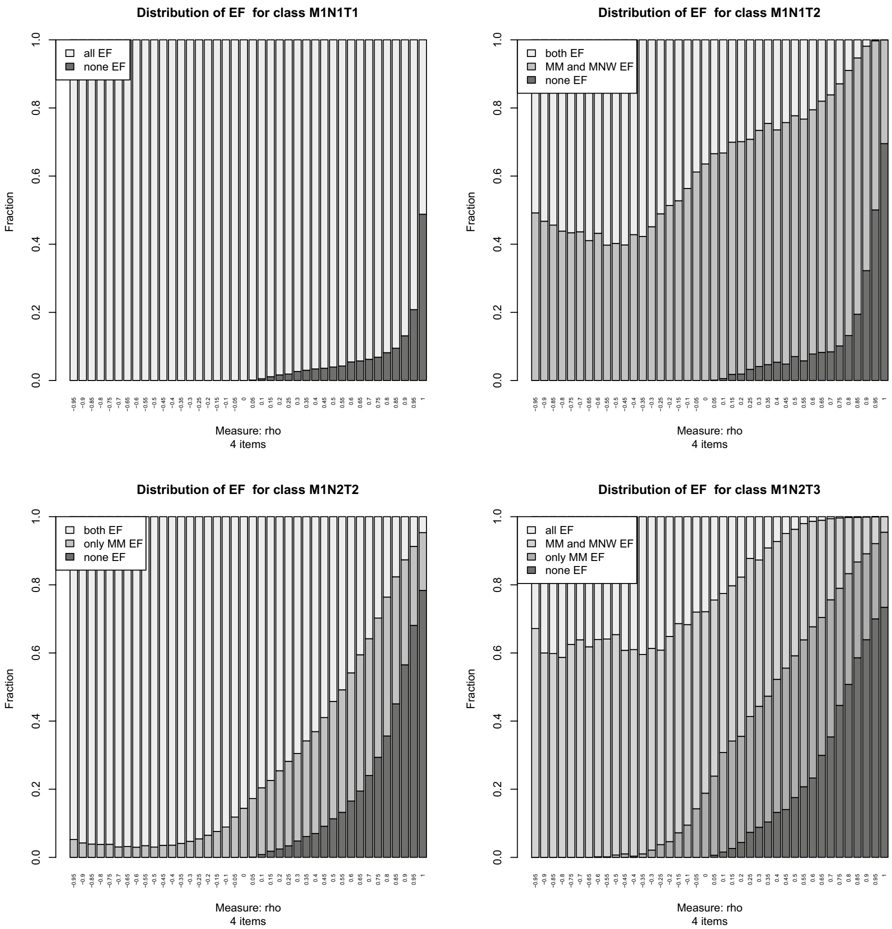

Fig. 15 Fractions of cases in which maximizing allocations are EF for different problem classes and levels of utility correlation: 4 items

items (in a problem where all three maximum allocations were different) and 0.007 for four items (in a problem where one allocation had all three properties).

Clearly, in a problem with two items, a negative correlation of utilities is a necessary and sufficient condition for an EF allocation to exist: If both players assign a higher utility to the same item, then no EF allocation can exist, and if their preferences differ, giving each its preferred item results in an EF allocation. Based on our simulation results, we therefore formulate the following

Conjecture 3 *For problems with an even number of items, or an odd number of items $n \geq 7$, a negative utility correlation is sufficient to guarantee that an EF allocation exists. For $n = 2$, this condition is also necessary.*

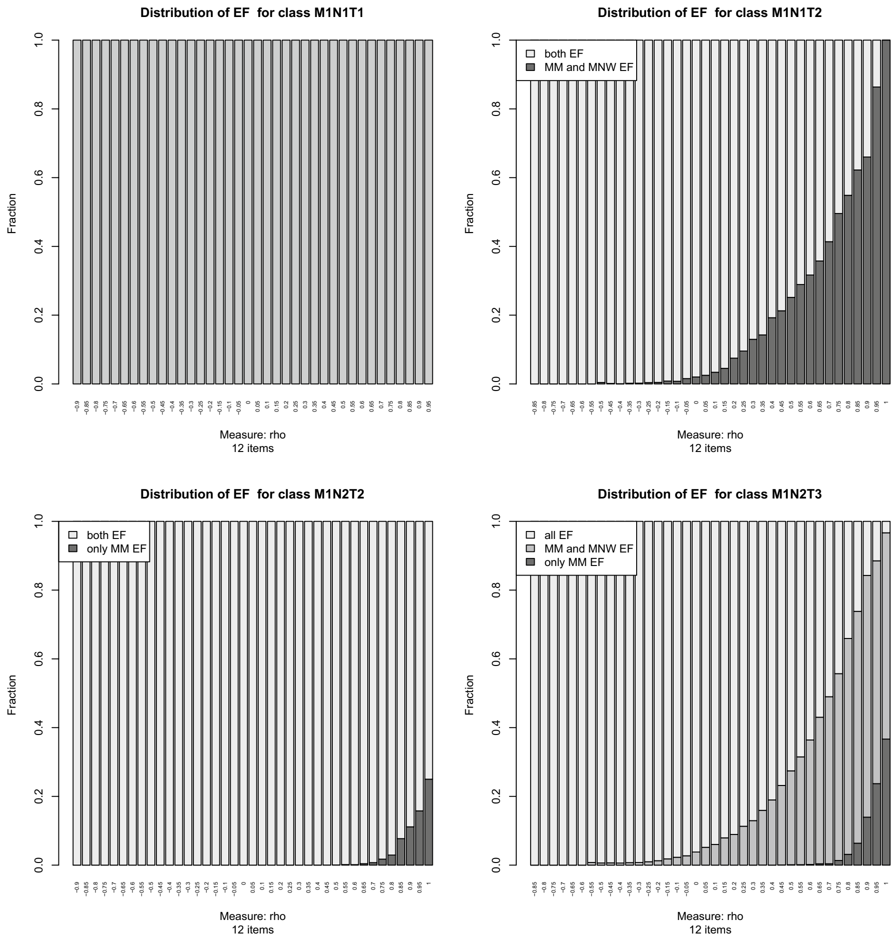

Fig. 16 Fractions of cases in which maximizing allocations are EF for different problem classes and levels of utility correlation: 12 items

Table 2 also shows that the threshold of correlations for S^{MM} not to be EF increases with problem size. Furthermore, in problems in class M1N1T1, where one allocation exhibits all three maximizing properties, only if utility correlation is very high is there no EF allocation. For larger problems, all problems in class M1N1T1 appear to have at least one EF allocation.

Table 2 Minimum correlations of utilities in problems where S^{MM}, S^{MNW}, and S^{MTW} allocations are or are not envy free

Items	Class	S^{MM}		S^{MNW}		S^{MTW}	
		Not EF	EF	Not EF	EF	Not EF	EF
4	M1N1T1	0.012	−1.000	0.012	−1.000	0.012	−1.000
	M1N1T2	0.026	−1.000	0.026	−1.000	−1.000	−1.000
	M1N2T2	0.012	−1.000	−1.000	−1.000	−1.000	−1.000
	M1N2T3	0.019	−1.000	−0.603	−1.000	−1.000	−0.998
5	M1N1T1	−0.935	−0.999	−0.935	−0.999	−0.935	−0.999
	M1N1T2	−0.937	−1.000	−0.937	−1.000	−0.995	−1.000
	M1N2T2	−0.890	−0.999	−0.989	−0.999	−0.989	−0.999
	M1N2T3	−0.844	−0.998	−0.980	−0.998	−0.998	−0.995
6	M1N1T1	0.516	−0.998	0.516	−0.998	0.516	−0.998
	M1N1T2	0.469	−0.994	0.469	−0.994	−0.978	−0.994
	M1N2T2	0.501	−0.997	−0.924	−0.997	−0.924	−0.997
	M1N2T3	0.533	−0.991	−0.904	−0.991	−0.991	−0.985
7	M1N1T1	0.951	−0.988	0.951	−0.988	0.951	−0.988
	M1N1T2	0.819	−0.994	0.819	−0.994	−0.945	−0.994
	M1N2T2	0.620	−0.988	−0.645	−0.988	−0.645	−0.988
	M1N2T3	0.591	−0.969	−0.688	−0.969	−0.934	−0.969
8	M1N1T1	0.982	−0.980	0.982	−0.980	0.982	−0.980
	M1N1T2	0.953	−0.984	0.953	−0.984	−0.942	−0.984
	M1N2T2	0.703	−0.992	−0.455	−0.992	−0.455	−0.992
	M1N2T3	0.728	−0.985	−0.202	−0.985	−0.917	−0.985
9	M1N1T1	–	−0.969	–	−0.969	–	−0.969
	M1N1T2	0.962	−0.961	0.962	−0.961	−0.804	−0.961
	M1N2T2	0.736	−0.949	−0.375	−0.949	−0.375	−0.949
	M1N2T3	0.787	−0.949	−0.156	−0.949	−0.851	−0.949
10	M1N1T1	–	−0.949	–	−0.949	–	−0.949
	M1N1T2	–	−0.970	–	−0.970	−0.640	−0.970
	M1N2T2	0.895	−0.921	−0.006	−0.921	−0.006	−0.921
	M1N2T3	0.817	−0.901	−0.004	−0.901	−0.887	−0.901
11	M1N1T1	–	−0.928	–	−0.928	–	−0.928
	M1N1T2	–	−0.893	–	−0.893	−0.660	−0.893
	M1N2T2	0.884	−0.938	0.030	−0.938	0.030	−0.938
	M1N2T3	0.863	−0.891	0.226	−0.891	−0.704	−0.891
12	M1N1T1	–	−0.910	–	−0.910	–	−0.910
	M1N1T2	–	−0.868	–	−0.868	−0.543	−0.868
	M1N2T2	–	−0.919	0.241	−0.919	0.241	−0.919
	M1N2T3	–	−0.874	0.256	−0.874	−0.580	−0.874

– no such allocations were found in the sample

8 Conclusions

The literature has identified a multitude of desirable properties that allocations of indivisible items should fulfil, properties that are not always compatible. We have studied these properties and their relationships in the specific setting of preferences that can be represented as additive cardinal utilities.

Like most collective decision problems, the fair division problems we have studied are characterized by a trade-off between fairness and efficiency (Alkan et al. 1991; Zukerman et al. 2008). In our setting, this trade-off is represented by the S^{MM}, S^{MNW}, and S^{MTW} allocations, for which we have identified a clear ordering. Allocation S^{MM} emphasizes equity, S^{MTW} emphasizes efficiency, while S^{MNW} balances these two properties. Propositions 2 (that S^{MNW} cannot be a different allocation if $S^{MM} = S^{MTW}$) and Proposition 3, about the ranking of utility differences, clearly reflect this ranking. This middle position of the Nash bargaining solution on the spectrum between equity and efficiency reflects the conclusions of Rachmilevitch (2015) in a somewhat different context.

Envy-freeness (EF) is an important property of allocations that reflects a concern for equity and fairness. Therefore, it is not surprising that among the three properties, MM exhibits the closest connection to EF. For instance, Brams et al. (2023) showed that if an EF allocation exists at all, then S^{MM} must be EF. Additional evidence for the ranking is our Corollary 3 that, if S^{MNW} is EF, then so is S^{MM}, and if S^{MTW} is EF, then S^{MM} and S^{MNW} must be too.

Our simulation analysis provides some quantitative insights into fairness and efficiency of allocations. It seems clear that, in practice, envy-freeness is an issue only if the number of objects to be allocated is small. For instance, we found that envy-free allocations exist in more than 99.9% of all problems with $n = 8$ items. Our millions of examples of problems with at least 11 items found at least one EF allocation in every one.

But while lack of equity (envy-free allocations) becomes a negligible problem when the number of items increases, the trade-off between equity and efficiency becomes more difficult. One reason is that the fraction of problems in which one allocation fulfils all desirable properties decreases rapidly. Most often, the maximin allocation differs from the other maximizing allocation(s). In contrast, the combined frequency of classes M1N1T1 and M1N2T2 (allocations with $S^{MNW} = S^{MTW}$, shown in figure 11) is surprisingly steady—it decreases from about 80 to almost 70% as the number of items increases 4–12.

Our simulation results also show that, besides the number of items, the level of conflict as measured by the correlation of the players' utilities has a strong influence on the properties of good allocations. The higher the correlation, the fewer the EF allocations, and the less likely that any one allocation can satisfy both equity and efficiency criteria.

In addition to our simulation results, we have offered three conjectures about the numbers of EF allocations and their relation to correlation. They offer new challenges to future research on fair division.

The first conjecture refers to the fraction of EF allocations in a problem. Figure 9 shows that the average fraction of allocations that are EF approaches 1/6 as problem size increases. That average is already close (around 15%) for the smallest problem size we tested (4 items). It should be noted that our conjecture concerns averages over all possible utility vectors—individual problems do have more or fewer EF allocations. A deeper analysis of this conjecture might ask how the expected fraction of EF allocations is related to the underlying probabilistic model of utilities, which might differ from our uniform distribution.

The second conjecture is that the maximum fraction of allocations that are EF is 1/4. As we already mentioned, it is easy to show that problems with this number of EF allocations exist, and that this conjecture fails if players value all items equally. Thus it seems that this upper bound depends on equal utilities being rare, so that there are items with utility exceeding $1/n$. Further analysis of cases in which this bound is reached is necessary to develop an understanding of when and how large numbers of EF allocations can arise.

Our last conjecture connected the existence of EF allocations to the correlation of utility values. Again, our observation may well depend on the uniform distribution that is the basis of our simulations. An attempt on an analytical proof of this conjecture will thus probably require assumptions about the joint distribution of two players' utilities.

Our observations have important consequences for the design of algorithms for fair division. For most problems, the search for an allocation that fulfils all desirable properties will fail. Any algorithm for fair division problems, at least one that applies when items are indivisible, must trade off between equity and efficiency. Our simulation results show that, as problem size increases, maximization allocations are more likely to be envy-free. If envy-freeness is sufficient for equity, then perhaps an allocation maximizing total value (MTW) is a good choice. Our results also show that this solution has a good chance of also being MNW, but it will frequently be different from the MM allocation.

Although our research has identified several relationships among properties of allocations in fair division problems, it is not without limitations. Our analytical results are based on additive cardinal evaluations, a specific model of preferences. In this setting, properties such as MM require that utilities be measured on the same scale, since otherwise the players' utilities would not be comparable. Additivity of utility functions also imposes restrictions on the generalizability of our results, ruling out the possibility of positive or negative synergies as well as saturation effects, which could play an important role in real-life fair division.

In addition to these general limitations, our simulation results exhibit the same limitations of generalizability as all computational studies. We performed a large number of simulations and took care to ensure that simulated preference values are uniformly distributed (in our case, over the n-dimensional unit simplex), but of course our results are limited to the cases that we analyzed. Cases not contained in our data might exist—unless they are ruled out by analytical results. For example, one million randomly generated problems with $n = 11$ and 12 did not uncover any examples with no EF allocation, but it is known that such examples exist when players' utilities are almost identical.

Funding Open access funding provided by University of Vienna. The authors did not receive specific funding for this study.

Declarations

Conflict of interest The authors have no conflict of interest to declare that are relevant to the content of this article.

References

Aleksandrov M, Ge C, Walsh T (2019) Fair division minimizing inequality. In: Progress in artificial intelligence: 19th EPIA conference on artificial intelligence, EPIA 2019, Vila Real, Portugal, September 3–6, 2019, proceedings, part ii 19, pp 593–605

Alkan A, Demange G, Gale D (1991) Fair allocation of indivisible goods and criteria of justice. Econometrica 59(4):1023–1039. https://doi.org/10.2307/2938172

Amanatidis G, Birmpas G, Filos-Ratsikas A, Hollender A, Voudouris AA (2021) Maximum Nash welfare and other stories about EFX. Theor Comput Sci 863:69–85. https://doi.org/10.1016/j.tcs.2021.02.020

Amanatidis G, Aziz H, Birmpas G, Filos-Ratsikas A, Li B, Moulin H, Voudouris AA, Wu X (2023) Fair division of indivisible goods: recent progress and open questions. Artif Intell 322:103965. https://doi.org/10.1016/j.artint.2023.103965

Barman S, Bhaskar U, Shah N (2020) Optimal bounds on the price of fairness for indivisible goods. In: Chen X, Gravin N, Hoefer M, Mehta R (eds) Web and internet economics. Springer International Publishing, Cham, pp 356–369

Bei X, Lu X, Manurangsi P, Suksompong W (2021) The price of fairness for indivisible goods. Theory Comput. Syst. 65(7):1069–1093. https://doi.org/10.1007/s00224-021-10039-8

Bentham J (1789) An introduction to the principles of morals and legislation. UCL Bentham Project, London

Bouveret S, Chevaleyre Y, Maudet N (2016) Fair allocation of indivisible goods. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of computational social choice. Cambridge University Press, Cambridge, pp 284–310

Brams SJ, Edelman PH, Fishburn PC (2003) Fair division of indivisible items. Theor Decis 55(2):147–180. https://doi.org/10.1023/B:THEO.0000024421.85722.0a

Brams SJ, Kilgour DM, Klamler C (2015) How to divide things fairly. Math Mag 88(5):338–348

Brams SJ, Kilgour DM, Klamler C (2017) Maximin envy-free division of indivisible items. Group Decis Negot 26(1):115–131. https://doi.org/10.1007/s10726-016-9502-x

Brams SJ, Kilgour DM, Klamler C, Wei F (2023) Two-person fair division of indivisible items: Bentham vs. Rawls on envy. J Philos 120(8):441–456

Butler J, Jia J, Dyer J (1997) Simulation techniques for the sensitivity analysis of multi-criteria decision models. Eur J Oper Res 103:531–546. https://doi.org/10.1016/S0377-2217(96)00307-4

Caragiannis I, Kaklamanis C, Kanellopoulos P, Kyropoulou M (2012) The efficiency of fair division. Theory Comput Syst 50(4):589–610. https://doi.org/10.1007/s00224-011-9359-y

Caragiannis I, Kurokawa D, Moulin H, Procaccia AD, Shah N, Wang J (2019) The unreasonable fairness of maximum Nash welfare. ACM Trans Econ Comput 7(3):12. https://doi.org/10.1145/3355902

Dickerson J, Goldman J, Karp J, Procaccia A, Sandholm T (2014) The computational rise and fall of fairness. In; Proceedings of the AAAI conference on artificial intelligence, vol 28(1). https://doi.org/10.1609/aaai.v28i1.8884

Freeman R, Sikdar S, Vaish R, Xia L (2019) Equitable allocations of indivisible goods. In: Kraus S (ed) IJCAI international joint conference on artificial intelligence, pp 280–286

Gamow G, Stern M (1958) Puzzle-math. Viking Press, New York

Goldman J, Procaccia AD (2015) Spliddit: unleashing fair division algorithms. SIGecom Exchanges 13(2):41–46. https://doi.org/10.1145/2728732.2728738

Halpern D, Procaccia AD, Psomas A, Shah N (2020) Fair division with binary valuations: one rule to rule them all. In: Chen X, Gravin N, Hoefer M, Mehta R (eds) Web and internet economics. Springer International Publishing, Cham, pp 370–383

Kilgour DM, Vetschera R (2018) Two-player fair division of indivisible items: comparison of algorithms. Eur J Oper Res 271(2):620–631. https://doi.org/10.1016/j.ejor.2018.05.057

Klamler C (2021) The notion of fair division in negotiations. In: Kilgour DM, Eden C (eds) Handbook of group decision and negotiation. Springer International Publishing, Cham, pp 81–109

Lipton RJ, Markakis E, Mossel E, Saberi A (2004) On approximately fair allocations of indivisible goods. In: Proceedings of the 5th ACM conference on electronic commerce. Association for Computing Machinery, New York, NY, USA, pp 125–131

Moulin H (2019) Fair division in the internet age. Annu Rev Econ 11:407–441

Nash JF (1950) The bargaining problem. Econometrica 18(2):155–162. https://doi.org/10.2307/1907266

Procaccia A (2016) Cake cutting algorithms. Handbook of computational social choice. Cambridge University Press, Cambridge, pp 311–330

Rachmilevitch S (2015) The Nash solution is more utilitarian than egalitarian. Theor Decis 79:463–478. https://doi.org/10.1007/s11238-014-9477-5

Rawls J (1971) A theory of justice. Belknap/Harvard University Press, Cambridge

Suksompong W (2023) A characterization of maximum Nash welfare for indivisible goods. Econ Lett 222:110956. https://doi.org/10.1016/j.econlet.2022.110956

Thomson W (2016) Introduction to the theory of fair allocation. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia A (eds) Handbook of computational social choice. Cambridge University Press, Cambridge, pp 261–283

Zukerman M, Mammadov M, Tan L, Ouveysi I, Andrew LL (2008) To be fair or efficient or a bit of both. Comput Oper Res 35:3787–3806. https://doi.org/10.1016/j.cor.2007.02.007