European Journal of Operational Research 271 (2018) 620–631

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Decision Support

# Two-player fair division of indivisible items: Comparison of algorithms

D. Marc Kilgour[a], Rudolf Vetschera[b,*]

[a] Wilfrid Laurier University Waterloo, Canada
[b] University of Vienna, Vienna, Austria

ARTICLE INFO

*Article history:*
Received 30 October 2017
Accepted 27 May 2018
Available online 1 June 2018

*Keywords:*
Fair division
Indivisible items
Borda score
Computational study

ABSTRACT

We study algorithms for allocating a set of indivisible items to two players who rank them differently. We compare eleven such algorithms, mostly taken from the literature, in a computational study, evaluating them according to fairness and efficiency criteria that are based on ordinal preferences as well as Borda counts. Our study is exhaustive in that, for every possible instance of up to twelve items, we compare the output of each algorithm to all possible allocations. We thus can search for "good" allocations that no algorithm finds. Overall, the algorithms do very well on ordinal properties but fall short on Borda properties. We also discuss the similarity of algorithms and suggest how they can be usefully combined.



## 1. Introduction

Fair allocation of resources is a central problem of collective decision-making. The study of fairness and efficiency criteria that reflect actors' preferences has motivated the development of algorithms to generate allocations that achieve these properties. It is not surprising that fair division problems have a pivotal position in the collective-choice literature (Klamler, 2010; Thomson, 2016; Young, 1994).

Fair division problems are difficult enough when resources are infinitely divisible, but become even more challenging when the resource to be allocated consists of a (finite) set of *indivisible* items. Each player is to receive a subset of the items, often called a "bundle"; allocations must be evaluated based on players' preferences over bundles. The simplest approach is to begin with a player's preference over individual items (e.g., a ranking) and "lift" (Bouveret, Chevaleyre, & Maudet, 2016) it to the level of bundles. To illustrate how difficult fair division problems can be, assume that utilities are linear (so that there are no synergies, positive or negative, among the objects), and consider the "Santa Claus problem" — to distribute a finite number of gifts to a finite number of children so as to maximize the utility of the unhappiest child. This problem is NP-hard, despite including only one fairness criterion. In fact, it cannot even be approximated efficiently (Bansal & Sviridenko, 2006).

Fair division of indivisible items is an important research topic, partly because sometimes no "fair" allocation exists, and partly because simple algorithms sometimes fail to find those that do. We treat fair division as a problem of social choice, not game theory; in other words, we assume that any algorithm can access true preference information, and do not account for individuals' incentives to conceal or distort that information. In fact, many social choice algorithms are available. On the one hand, there are plausible, common-sense approaches, such as asking players, in sequence, to claim their most preferred unallocated item, which mimics the "I cut, you choose" cake-cutting procedure. On the other hand, the recent academic literature proposes many algorithms that specifically aim to find allocations satisfying various fairness and efficiency criteria (e.g., Brams, Kilgour, & Klamler, 2012; 2015; 2017; Darmann & Klamler, 2016; Pruhs & Woeginger, 2012).

The aim of this paper is to compare fair-division algorithms for indivisible items. We draw from across the literature, and adopt its most common restrictions, considering only algorithms that produce balanced allocations to two individuals. An allocation is *balanced* if each individual receives the same number of items. Balanced allocations satisfy a very weak notion of fairness, one that can be applied without preference information. The two-player setting is a sub-problem with practical implications, for example the division of marital property in a divorce. Good algorithms for two-person balanced allocations may also provide important clues about how to approach more general problems.

We further restrict our focus to algorithms that consider each item only once, immediately assigning it to one player or the other. We do not include algorithms that leave a "Contested Pile" of unallocated items to be distributed in a second stage, perhaps by a different algorithm. Our algorithms range from "naive" methods to recent proposals in the academic literature. We compare algorithms by applying various criteria measuring efficiency and fairness to

* Corresponding author.
*E-mail addresses:* mkilgour@wlu.ca (D.M. Kilgour), rudolf.vetschera@univie.ac.at (R. Vetschera).

https://doi.org/10.1016/j.ejor.2018.05.057

the allocations they generate. All of our criteria are based on ordinal preferences, including some criteria based on Borda scores, which can be thought of as a crude approximation to cardinal utilities (Darmann & Klamler, 2016).

Our computational study is, to the best of our knowledge, the first of its kind. Randomized computational studies have been carried out in other fields, for example to assess voting rules (Laslier, 2010; Merrill, 1988). Our approach is not only to compare algorithms to each other, but also to assess how well they do as a group. For each problem size, we consider all possible profiles of individual rankings, enabling us to make exact statements about how often an event occurs, rather than the probabilistic statements that would follow from a sampling approach. For each problem, we consider all allocations that each algorithm might find. (All of our algorithms require at least one arbitrary choice, such as which player acts first.) We assess an algorithm on the quality of all of the allocations it produces, keeping track of the frequency of branching and the number of allocations generated on average and in extreme cases. We see these characteristics of an algorithm as highly relevant in practical problems. Also, for each problem, we find all possible balanced allocations; knowing all "good" allocations, we can determine how many each algorithm finds, and whether there are interesting allocations not found by any algorithm.

In order to analyze all possible preference profiles for each problem size, we restricted our computational study to problems of up to 12 items. Although this size restriction is quite severe, we believe that the analysis of small problems is valuable. First, we are able to demonstrate clear differences among algorithms even within this range. Second, in smaller problems, the link from algorithm characteristics to effects often becomes clear. Third, the effects of indivisibility seem stronger when the number of items is small, suggesting that the interesting features of algorithms are most likely to be manifest in small problems. Finally, many real-world problems (such as the allocation of cabinet seats to political parties) involve only a small number of items.

The paper proceeds with a discussion of the fairness and efficiency properties of allocations used in comparing algorithms (Section 2), followed by a description of the algorithms and their implementation (Section 3). Then (Section 4) we elaborate on exactly how we compare algorithms, and formulate the detailed research questions we want to study. After presenting our results in detail (Section 5) we draw some conclusions (Section 6), including some ideas about important questions that remain unanswered.

## 2. Properties of allocations

We consider a set $S$ of $|S| = N$ items to be allocated to players $A$ and $B$. Let $M = \{A, B\}$. For a player $m \in M$, denote the opponent of $m$ by $\overline{m}$. An allocation $\mathbf{X} = (X_A, X_B)$ assigns subsets $X_A \subseteq S$ to $A$ and $X_B \subseteq S$ to $B$, provided $X_A \cap X_B = \emptyset$ and $X_A \cup X_B = S$. We assume throughout that $N$ is even and consider only allocations $\mathbf{X} = (X_A, X_B)$ that are *balanced* in that $|X_A| = |X_B| = N/2$. Denote the set of all balanced allocations of $S$ by $\mathcal{A}$.

To evaluate allocations, one needs a model of players' preferences over sets of items. We focus on preferences on subsets of $S$ that can be obtained from (strict) rankings of $S$. Denote player $m$'s ranking of items by $\succ_m$. We construct a partial ordering of subsets of $S$ proposed by Brams et al. (2012). Let $X \subseteq S$ and $Y \subseteq S$ satisfy $|X| = |Y|$. Then $X$ is *ordinally less* than $Y$ for player $m$, denoted by $X \prec_m^o Y$, if there exists an injective mapping $f$: $Y \to X$ so that $\forall y \in Y$: $y \succ_m f(y)$. If $X \prec_m^o Y$, then $Y \succ_m^o X$, and we say that $Y$ is *ordinally more* than $X$. Based on these relations, we define several criteria for (balanced) allocations.

Pareto optimality is an efficiency property. An allocation $\mathbf{X}$ is *Pareto optimal (PO)* if there is no other allocation $\mathbf{Y}$ such that, for both players, the set of items received under $\mathbf{X}$ is ordinally less than the set of items received under $\mathbf{Y}$, i.e. allocation $\mathbf{X}$ is *PO* iff

$$\nexists \mathbf{Y} \in \mathcal{A} : \forall m \in M : Y_m \succ_m^o X_m \tag{1}$$

Note that we define Pareto optimality as the absence of strict dominance. We simplify the usual treatment that allows some player to be indifferent between its two allocations because underlying orderings are strict; thus, if all items are allocated, any change in the allocation to one player also changes the allocation to the other, making subset indifference impossible.

The second property we consider is *envy-freeness (EF)*. An allocation $\mathbf{X}$ is envy-free if each player prefers its set to the complement, the set allocated to the opponent (Klamler, 2010). An allocation $\mathbf{X}$ is *EF* iff

$$\forall m \in M : X_m \succ_m^o X_{\overline{m}} \tag{2}$$

A third criterion of fairness for allocations is the max-min property. An allocation is *max-min (MM)* iff the maximum rank of any item allocated to a player in that player's ranking is minimal (Brams, Kilgour, & Klamler, 2017). Denote by $r_m(i) = |\{x \in S : x \succ i\}| + 1$ the rank of item $i$ for player $m$. The most preferred item has rank 1 and the least preferred item rank $N$. An allocation $\mathbf{X}$ is *MM* iff

$$\max_{m \in M} \max_{i \in X_m} r_m(i) = \min_{\mathbf{Y} \in \mathcal{A}} \max_{m \in M} \max_{i \in Y_m} r_m(i) \tag{3}$$

The three properties just defined are commonly used in the literature on fair division of indivisible items—in fact, many of the algorithms we study were specifically designed to fulfill some of them—but they are not without drawbacks. Most notably, ordinally less is not a complete relation on the set of all (balanced) allocations. For any property defined as the non-existence of a related allocation (e.g., an allocation is Pareto optimal if there exists no other allocation providing preferred subsets to both players), one could argue that a complete relation would be more demanding because it offers more potential violations. In other words, as a basis for assessment, an incomplete relation is rather weak.

We therefore complement properties based on the *ordinally less* relation with properties based on Borda scores. Recall that player $m$'s ranking of item $i \in S$ is $r_m(i)$. The Borda score of a set $X_m \subseteq S$ for player $m$ is

$$B_m(X_m) = \sum_{i \in X_m} (N + 1 - r_m(i)). \tag{4}$$

When only ordinal information is available, Borda scores can be regarded as an approximation to cardinal utilities (Darmann and Klamler, 2016, p. 545), and provide a complete preorder on subsets. The ordinally less relation obviously implies the corresponding ranking of Borda scores.

Based on Borda scores, we define four properties of allocations that are analogous to properties based on the ordinally less relation. The first property transfers Pareto optimality to Borda scores. A balanced allocation $\mathbf{X}$ is *Borda Pareto optimal (BP)* iff

$$\nexists \mathbf{Y} \in \mathcal{A} : \forall m \in M : B_m(Y_m) \geq B_m(X_m) \wedge \exists n \in M : B_n(Y_n) > B_n(X_n) \tag{5}$$

Using Borda scores, we can also apply the utilitarian concept of efficiency as maximization of the sum of scores (Bertsimas, Farias, & Trichakis, 2012). A balanced allocation $\mathbf{X}$ satisfies *maximal Borda sum (BS)* iff

$$\sum_{m \in M} B_m(X_m) = \max_{\mathbf{Y} \in \mathcal{A}} \sum_{m \in M} B_m(Y_m) \tag{6}$$

Because *BS* clearly implies *BP*, we concentrate on the *BS* property.

When cardinal evaluations are available, the concept of envy-freeness can be applied by requiring that no player assigns a