

How to Allocate Hard Candies Fairly

Marco Dall'Aglio and Raffaele Mosca
Dipartimento di Scienze
Università d'Annunzio
Viale Pindaro, 42
65127 — Pescara, Italy

Abstract

We consider the problem of allocating a finite number of indivisible items to two players with additive utilities. We design a procedure that looks for all the maximin allocations. The procedure makes repeated use of an extension of the Adjusted Winner, an effective procedure that deals with divisible items, to find new candidate solutions, and to suggest which items should be assigned to the players.

JEL classification: D61, D63

Keywords: Fair division, indivisible items, Adjusted Winner

1 Introduction

This paper presents a procedure for allocating a set of indivisible items between two players with subjective preferences over the items.

While most of the literature in fair division theory deals with one or more completely divisible goods (such as cakes or pieces of land), recent works by Brams with several coauthors drew attention on the problem of allocating several indivisible items. Brams, Edelman and Fishburn [4] point out how the most commonly accepted criteria for optimality may conflict with each other when players rank the items according to their preferences, so that

achieving fairness requires some consensus on the ground rules and some delicacy in applying them.

In the same context of ordinal preferences, Brams and King [8] focus on the possible incompatibility between rank- and Borda-maxmin allocations on one side and envy-free ones on the other. Brams, Edelman and Fishburn [5], again, provide conditions for the existence of allocations which are optimal according to different criteria and study the relationship among those criteria for any number of players and items. An earlier work by Brams and Fishburn [6], focuses on the case of two players with the same ranking on the items.

When it comes to the design of specific procedures, however, it turns out that most of the proposals devise some technique to treat some, or all, of the contended items as divisible. This is the case, for instance, of the *Adjusted Winner* (AW) procedure, certainly the most popular and effective procedure so far conceived. In case no actual splitting is allowed, one may recur to such

surrogates as the use of side payments, as in the Knaster method, or that of randomization, where items are given according to a probability distribution, or of side payments to compensate the giving up of some item. As noted in [13], however, there are situations where these methods are impractical or impossible to implement.

If we focus on methods that deal exclusively with the allocation of indivisible items, with no side actions to mitigate the discontent of some players, we find, quite surprisingly, a more narrow choice. Classical methods for the 2-player case are described in [9] and, especially, [10]. The simplest method is that of *strict alternation* of the players' right to pick an item. This approach usually favors the player who picks first. To mitigate this advantage, Brams and Taylor propose a preliminary ranking of the items (a "query step") with the immediate assignment of the undisputed items, followed by a division of the remaining items (the "contested pile") via an alternation schemes that uses a more balanced sequence. The whole scheme is called *balanced alternation*. Another option is given by Lucas' *method of markers* which could be seen as a discrete variant of the sliding knife procedure.

A few recent additions complete the list. Herreiner and Puppe in [13] define a *descending demand* procedure where each player, in turn, declare their most preferred bundle (i.e. a collection of items) until a feasible arrangement is met that maximizes the bundle's rank of the least favored player. The findings in [6] suggest a procedure that Brams and Fishburn set out in the same work to single out an allocation which is Pareto-optimal, it ensures that the less well-off player does as well as possible, and, often, he/she does not envy the other player. In a similar fashion Brams and King [8] devise a simple procedure based on balanced alternation and sincere choices that yields Pareto-optimality and does not rule out envy-freeness¹

All the above mentioned methods require the players' ability to rank items or bundles of them, and can be adapted to the simpler framework in which players are able to assess the subjective utility (or score) of each item and these evaluations are additive. Brams and Fishburn [6] show conditions that make preference relations compatible with additive utilities, and explain how to simplify their procedure in this situation. Anyway, we record the lack of a procedures specifically designed to work with additive utilities, in a manner similar to what the AW procedure does for the divisible case. Our aim indeed is to devise a specific procedure that makes repeated use of the original AW procedure as a guide to decide who gets the single items. The procedure mimics the branch-and-bound algorithms of Operations Research (OR), but keeps the procedural appeal of the original AW and it can be implemented as a simple set of instructions given to the players. The association of ideas from OR with fair division is not new: In [15] Kuhn defines a linear program that has the Knaster rule for the efficient allocation of items with side payments as its solution. Demko and Hill [12] define a maximin optimization problem. They show

¹In the context of ordinal preferences allocations are divided into *envy-free*, *envy-possible* and *envy-ensuring* ones. The procedure returns an allocation belonging to the first two classes

that this problem is computationally intractable and provide a lower bound for optimal value. The second half of the paper deals with randomized solutions for the same problem and shows how these can be computed through linear programming and duality techniques.

We adopt the same framework, focusing on the case of two players. Each player assigns a non-negative value to each item. The evaluations are additive, but no normalization is required, so the total value of the items may differ for the two players.

This work does not deal with manipulability issues: is it advantageous for the players to reveal the items' true values? A discussion of the manipulability for the Adjusted Winner procedure appears in [9] and strategy proof procedures have recently been introduced in [7] for the divisible case.

A similar approach to the one presented here is being developed by Bezáková and Dani [3]. The purpose of the two works, however, differs. In [3] computationally efficient algorithms are given that approximate the optimal solution and are implemented by suitably programmed computer routines. Here we focus on exact solutions, with the main aim to extend the AW procedure to cover the case of indivisible items and keep its procedural nature.

Section 2 defines the problem. Section 3 takes another look at the AW algorithm and an extension is discussed to consider the situations where players own initial endowments. Incidentally, a more efficient version of the original procedure is considered for the case where a large number of items are at stake. Finally, section 4 illustrates the branch-and-bound algorithm that makes use of the AW procedure with initial endowments to find new candidate solutions, and to suggest which items should be forcedly assigned to the players.

2 The problem

We consider the following simple problem: two children (players), Alice and Bob, are given a set of m hard candies to be shared between themselves. Candies are indivisible and each of them is assigned to one of the children. Children value the sweets according to their own taste. An allocation is sought that is optimal according to some social welfare criterion.

More formally, let $M = \{1, \dots, m\}$ be the set of disputed items and let a_1, a_2, \dots, a_m (b_1, b_2, \dots, b_m resp.) be the non-negative evaluations of the single items by Alice (Bob, resp.). An *integer allocation* for the m items is described by a vector $x = (x_1, \dots, x_m) \in \{0, 1\}^m$. If $x_i = 1$ (resp. $x_i = 0$), then item i goes to Alice (Bob, resp.). The satisfaction (or score) of the two players is given by, respectively,

$$v_A(x) = \sum_{i \in M} a_i x_i \quad \text{and} \quad v_B(x) = \sum_{i \in M} b_i (1 - x_i) \quad (1)$$

There are many criteria that mediate between the conflicting interests of the players. We follow Brams and Fishburn [6], who

recommend an alternative procedure that implements [their] fairness criteria when additive utilities are presumed. [...]the alternative procedure seeks a division that maximizes $[\min\{v_A, v_B\}]$ over all divisions, subject to

$$v_A((1, 1, \dots, 1)) = v_B((0, 0, \dots, 0)) \quad (2)$$

Therefore we look for an integer allocation that achieves

$$z^* = \max \{ \min\{v_A(x), v_B(x)\} : x \in \{0, 1\}^m \} \quad (\text{IFD})$$

subject to (2). The same problem was considered earlier by Demko and Hill in [12], who noted its NP-hardness. In fact, assume that $a_i = b_i$ for every $i \in M$: then solving (IFD) gives an answer to the problem of finding a partition of a set of positive integers in two subsets of equal sum, which is NP-complete (see for instance [19]).

As pointed out by Brams, Edelman and Fishburn [4] in the context of ordinal preferences, a maximin allocation of indivisible items may generate envy between the players. Moreover the optimal partition may assign a different number of items to the players — thus being *unequal*. *Equitability*, i.e. the property that the scores of the two players coincide, is rarely obtained for a solution of (IFD). Moreover, the solution may not be unique. Quoting Brams and Fishburn again

If two or more division maximize the min value, [the procedure] then finds an [allocation] within the maximin set that maximizes $\max\{v_A(x), v_B(x)\}$

This is referred to in the literature as the *equimax* (or Rawls, or Dubins-Spanier) allocation and has the property of (strong) Pareto-optimality: no other allocation weakly dominates it. Alternatively, we may choose a maximin allocation that minimizes $\max\{v_A(x), v_B(x)\}$. This time only weak Pareto-optimality is ensured (no other allocation strongly dominates it) but the resulting allocation would be closer to equitability. Neither one of the two restrictions, however, would ensure uniqueness in the solution. In what follows, we will distinguish between methods that are able to find one maximin solution, and those who can list them all.

3 Relaxing Integer Fair Division: The Adjusted Winner procedure

Suppose now that children are given muffins (with different flavors), instead of hard candies. Each muffin can be given in its entirety to one of the children — or it can be split in any proportion. We are now dealing with the allocation of m divisible items between two players.

It is further assumed that all items $i \in M$ are homogeneous. Thus player 1 can receive a part $x_i \in [0, 1]$ of item i , while player 2 gets the rest. The two players will benefit, respectively, by $x_i a_i$ and $(1 - x_i) b_i$ from the splitting. The overall satisfaction of each player is still given by (1). We now look for an allocation $x = (x_1, x_2, \dots, x_m) \in [0, 1]^m$ that achieves

$$z^+ = \max \{ \min \{ v_1(x), v_2(x) \} : x \in [0, 1]^m \} \quad (\text{DFD})$$

with v_A and v_B satisfying (2). As noted in [12], (DFD) can be solved through linear programming, and in the OR jargon, this is the *linear relaxation* of (IFD).

Here we are going to show that a solution for (DFD) is readily provided by a popular and effective step-by-step procedure in fair division, known as *Adjusted Winner*.

3.1 The Adjusted Winner algorithm

The Adjusted Winner (AW) algorithm was introduced by Brams and Taylor in [9] (with many applications analyzed in [10]). Their aim was to provide a step-by-step procedure returning a partition that is equitable, Pareto optimal and envy-free (in the sense that none of the player feels that the other player has received more than him/herself). A brief sketch of the algorithm follows — for a more detailed account we refer to [9] and [10]. There are two phases:

the “winning” phase. each player temporarily receives the items that he/she values more than the other player does — ties being temporarily assigned to any of the players. The total score of each player, v_A and v_B respectively, is computed.

the “adjusting” phase. Items are transferred, one at a time from the “richer” player to the “poorer” one, starting with the items with ratio a_i/b_i closer to 1. To reach equitability one item may be split into two parts.

As an exemplification, suppose $v_A \geq v_B$. Then Alice begins transferring items to Bob, one at a time, starting with the item with ratio a_i/b_i closer to 1 (and greater than or equal to 1). The handover continues until perfect equitability is achieved, or the roles of the “richer” and “poorer” player are reversed. In the last case, suppose that after the handover of, say, item r we have $v_A < v_B$. Item r is then split, with Alice getting a fraction given by

$$x_r = \frac{b_r + v_B^{-r} - v_A^{-r}}{a_r + b_r}$$

where v_A^{-r} and v_B^{-r} are the scores obtained by the two players so far in the process *without* considering item r . Bob gets the remaining fraction. The item is split according to the same proportions also when Bob is favored

in the “winning” phase and the handover occurs in the opposite direction. Both players walk out of the procedure with a common score of

$$z^+ = v_A = v_B = \frac{v_B^{-r} a_r + v_A^{-r} b_r + a_r b_r}{a_r + b_r}$$

We next show that AW provides exactly what we are looking for in the maximin problem with divisible items.

Proposition 3.1. *The AW algorithm solves (DFD). Therefore, the AW solution is also maximin.*

Some preliminary results are required. First of all consider the *allocation range*.

$$\mathcal{D} = \{(v_A(x), v_B(x)) : x \in [0, 1]^m\}$$

Lemma 3.2. *\mathcal{D} is a convex and compact set in \mathbb{R}^2 .*

Proof. Pick $x, y \in [0, 1]^m$ and $\gamma \in [0, 1]$. Then

$$v_A(\gamma x + (1 - \gamma)y) = \gamma v_A(x) + (1 - \gamma)v_A(y)$$

and the same holds for v_B , so \mathcal{D} is convex. Compactness is a consequence of the compactness of $[0, 1]^m$ and the continuity of v_A and v_B . More in detail, $\mathcal{D} \subset [0, \sum_{i \in M} a_i] \times [0, \sum_{i \in M} b_i]$, so \mathcal{D} is bounded. Consider now a sequence $\{x_n\}$ in $[0, 1]^m$ for which $(v_1(x_n), v_2(x_n))$ converges. Since $[0, 1]^m$ is compact, there exists a subsequence $\{x_{n'}\}$ converging to some $x^* \in [0, 1]^m$. Since v_A and v_B are continuous, we have

$$(v_A(x_{n'}), v_B(x_{n'})) \rightarrow (v_A(x^*), v_B(x^*)) \in \mathcal{D}$$

and \mathcal{D} is closed. □

Next we characterize the maximin solutions.

Lemma 3.3. *A maximin solution always exists. An allocation is maximin if and only if it is Pareto optimal and equitable.*

Proof. We prove the “only if” part of the prove, since this is what is actually needed for Proposition 3.1.²

We consider the set \mathcal{D} of all the allocations’ values. An allocation x is Pareto if there is no other point of \mathcal{D} in the upper quadrant pointed on $(v_A(x), v_B(x))$ (with the exception of x itself). The allocation is equitable if $(v_A(x), v_B(x))$ lies on the bisector of the positive quadrant.

Let \mathcal{Q} be the family of upper quadrants pointed on the equitable allocations. A maximin solution is obtained by considering the supremum of the quadrants in \mathcal{Q} that intersects \mathcal{D} . Since \mathcal{D} is compact, the supremum is attained, and a maximin solution x^* exists.

²For the whole proof we refer to the longer version of the paper.

Suppose that x is Pareto and equitable. Equitability implies that the upper quadrant pointed on $(v_A(x), v_B(x))$ is in \mathcal{Q} . Pareto optimality implies that no other part of \mathcal{D} lies on the same quadrant. Therefore x is maximin. \square

Proof of Proposition 3.1. It is a straightforward consequence of Lemma 3.3 and the following

Theorem 3.4. (*Brams and Taylor, [9], Th.4.1*) *AW produces an allocation of the goods that is Pareto-optimal and equitable.*

\square

An alternative proof of Theorem 3.4 which links AW to a cake cutting scheme is provided by Jones [14].

3.2 The maximin problem with initial endowments

The AW procedure is flexible enough to cover the situation where the two players own initial endowments. This variation is interesting in its own rights. An optimal allocation is sought when the utility of each player is the sum of the initial endowment and the values of the items (or fractions thereof) received. Our interest in this problem, however, is mainly instrumental. In order to implement a branch-and-bound method for the case of indivisible items we need to solve several instances of the corresponding problem with divisible items in which certain items are forcedly assigned to the players. These items represent their initial wealth. Let $\alpha \geq 0$ (resp. $\beta \geq 0$) the initial endowment of Alice (Bob, resp.) that add up to the players' utilities.

The problem of interest is now:

$$z^+ = \max \{ \min\{\alpha + v_A(x), \beta + v_B(x)\} : x \in [0, 1]^m \} \quad (\text{DFD-ie})$$

This time we do not impose a normalization condition such as (2), but rather assume all terms a_i, b_i to be strictly positive. We could in fact set up a preliminary step that deals with null values. If $a_i > 0$ and $b_i = 0$, then item i can be assigned to Alice with no harm for Bob, and increase her initial endowment. Similarly, Bob could take all items with null value to Alice (and items with no value for both could be thrown away). It may be worthwhile noticing that, by removing (2), we may lose envy-freeness. Consider for instance the case with no initial endowments and only one muffin M , with $v_A(M) = 10$ and $v_B(M) = 5$. Bob will now get 2/3 of the muffin, leaving Alice envious.

Once again the maximin solution coincides with the Pareto and equitable solution, but only when the value of the assignable items according to the poorer player is larger than or equal to the difference between the initial endowments. We propose the following:

The Adjusted Winner procedure with initial endowments (AW-ie)

Case 1 If $\sum_{i \in M} b_i \leq \alpha - \beta$ assign all the items to Bob. The maximin value will be $z^+ = \beta + \sum_{i \in M} b_i$

Case 2 If $\sum_{i \in M} a_i \leq \beta - \alpha$ then assign all the items to Alice and $z^+ = \alpha + \sum_{i \in M} a_i$

Case 3 If $-\sum_{i \in M} a_i < \alpha - \beta < \sum_{i \in M} b_i$ then start an AW procedure with the only difference that the initial endowment is taken into account to reach an equitable allocation. So, after a winning phase identical to the AW procedure, the total scores, inclusive of the initial endowment are computed. In the adjusting phase, items are transferred from the “richer” player to the “poorer” one, ordering the items in terms of the preference ratios a_i/b_i . The process stops when perfect equitability between the scores with the endowments is reached or the roles of the two players are reversed. In this case, the last transferred item, say r , is split and Alice gets a fraction given by

$$x_r = \frac{\beta - \alpha + b_r + v_B^{-r} - v_A^{-r}}{a_r + b_r}$$

while Bob gets the rest. Both players walk out with a common score of

$$z^+ = \alpha + v_A = \beta + v_B = \frac{(\beta + v_B^{-r})a_r + (\alpha + v_A^{-r})b_r + a_r b_r}{a_r + b_r}$$

Proposition 3.5. *The AW-ie procedure returns the solution for (DFD-ie).*

Proof. In this case the utility of the two players is given, respectively, by $v'_A(x) = \alpha + v_A(x)$ and $v'_B(x) = \beta + v_B(x)$, so the new allocation range \mathcal{D}' is simply the allocation range \mathcal{D} translated by (α, β) . The three cases listed above correspond to different positions of \mathcal{D}' with respect to the bisector of the first quadrant

Case 1: \mathcal{D}' lies below the bisector. So $u \geq w$ for all $(u, w) \in \mathcal{D}'$ and the allocation that assigns all goods to Bob is maximin.

Case 2: \mathcal{D}' lies above the bisector. Therefore $u \leq w$ for all $(u, w) \in \mathcal{D}'$ and all goods are given to Alice in order to have the highest maximin value.

Case 3: \mathcal{D}' crosses the bisector Lemma 3.3 is still valid. Thus, the procedure looks for an allocation that is Pareto-optimal and equitable.

□

If case 3 holds, a more efficient version of AW can be implemented just as it was done for the original AW. This time $v_A^i(\lambda)$ and $v_B^i(\lambda)$, $i = 1, 2$ will include the initial endowments of the respective players.

4 A branch and bound algorithm

When solving the maximin allocation problem (IFD) there is a finite number of possible candidates to choose from. In principle the solution can be obtained in finite time by computing the value of each allocation for the two players. This process can be considerably speeded up if we consider a branch-and-bound technique that splits the original problem into smaller subproblems and uses upper bounds to avoid exploring certain parts of the set of feasible integer solutions. This approach makes repeated use of the Adjusted Winner procedure with initial endowment and keeps the procedural character of the latter.

In what follows, we will consider a series of constrained subproblems in which some of the items have already been assigned to the players. Let $A, B \subset M$, with $A \cap B = \emptyset$. Let $S(A, B)$ be the constrained problem in which the items in A (B , resp.) are assigned to Alice (Bob, resp.), i.e., $x_i = 1$ for each $i \in A$ ($x_i = 0$ for each $i \in B$). $S(\emptyset, \emptyset)$ denotes the original (unconstrained) problem.

For a given couple of disjoint index sets, A, B in M , let $\bar{x}(A, B)$ denote a feasible allocation for the constrained problem and let $\bar{z}(A, B)$ denote the corresponding value. Moreover, let $x^*(A, B)$ and $z^*(A, B)$ denote the solution and the value of $S(A, B)$. Finally let $x^+(A, B)$ and $z^+(A, B)$ be, respectively, the solution and value for the linear relaxation of $S(A, B)$, i.e. for the case where splitting of the contended items is allowed. Clearly, the following holds for each couple of A and B :

$$\bar{z}(A, B) \leq z^*(A, B) \leq z^+(A, B) \quad (3)$$

The results in Section 3 can be used to compute $x^+(A, B)$ and $z^+(A, B)$. In particular we set $\alpha = \sum_{i \in A} v_1(x_i)$ and $\beta = \sum_{i \in B} v_2(x_i)$, and divide the remaining $M' = M \setminus (A \cup B)$ items according to the AW-ie procedure. Since $x^+(A, B)$ contains at most one fractional component, $\bar{x}(A, B)$ may be obtained by approximating the fractional coordinate to the nearest integer, 0 or 1.

4.1 A variable elimination test

The branch-and-bound procedure defines a series of subproblems in which an increasing numbers are forcedly assigned to one player or the other. Since the procedure becomes simpler as the number of pre-assigned items increases, and following [18], p.452, we consider a variable elimination test that, for any given subproblem, checks whether additional items can be assigned priori to any further analysis.

Let $A, B \subset M$ be a couple of disjoint sets of items and take $i \in M'$.

Proposition 4.1. (a) *If*

$$z^+(A \cup \{i\}, B) < \bar{z}(A, B) \quad (4)$$

then $x^(A, B \cup \{i\})$ solves $S(A, B)$, while $x^*(A \cup \{i\}, B)$ does not.*

(b) If

$$z^+(A, B \cup \{i\}) < \bar{z}(A, B) \quad (5)$$

then $x^*(A \cup \{i\}, B)$ solves $S(A, B)$, while $x^*(A, B \cup \{i\})$ does not.

Proof. By assumption and (3) we have

$$z^*(A \cup \{i\}, B) \leq z^+(A \cup \{i\}, B) < \bar{z}(A, B) \leq z^*(A, B)$$

So $x^*(A \cup \{i\}, B)$ cannot be a solution for $S(A, B)$. If this is the case, then $x^*(A, B \cup \{i\})$ must be a solution for the same problem. Part (b) is established symmetrically. \square

The result simply states that whenever condition (4) ((5), resp.) occurs, then $S(A, B)$ can be replaced by $S(A, B \cup \{i\})$ ($S(A \cup \{i\}, B)$, resp.). When the two sides of (4), or (5), attain equality, there is a partial extension of the previous result:

Proposition 4.2. (a) If $z^+(A \cup \{i\}, B) \leq \bar{z}(A, B)$, then either $x^*(A, B \cup \{i\})$ or $\bar{x}(A, B)$ solve $S(A, B)$.

(b) If $z^+(A, B \cup \{i\}) \leq \bar{z}(A, B)$, then either $x^*(A \cup \{i\}, B)$ or $\bar{x}(A, B)$ solve $S(A, B)$.

(c) If $z^+(A \cup \{i\}, B) \leq \bar{z}(A, B)$ and $z^+(A, B \cup \{i\}) \leq \bar{z}(A, B)$, then $\bar{x}(A, B)$ solves $S(A, B)$.

Proof. (a) By assumption

$$z^+(A \cup \{i\}, B) \leq \bar{z}(A, B) \leq z^*(A, B)$$

Assume now that $x^*(A, B \cup \{i\})$ does not solve $S(A, B)$. Then $x^*(A \cup \{i\}, B)$ will work instead, and thus

$$z^*(A, B) \leq z^*(A \cup \{i\}, B) \leq z^+(A \cup \{i\}, B)$$

Comparing the two inequalities, we conclude that $\bar{z}(A, B) = z^*(A, B)$ and $\bar{x}(A, B)$ solves $S(A, B)$. Part (b) is proved with a symmetrical argument.

(c) By definition

$$\bar{z}(A, B) \leq z^*(A, B) \leq z^+(A, B) \leq \max\{z^+(A \cup \{i\}, B), z^+(A, B \cup \{i\})\}$$

while the hypotheses reads

$$\max\{z^+(A \cup \{i\}, B), z^+(A, B \cup \{i\})\} \leq \bar{z}(A, B)$$

Thus $\bar{x}(A, B)$ solves $S(A, B)$. \square

The use of Proposition 4.2 is more subtle: when situation (a) occurs, than we replace $S(A, B)$ with $S(A, B \cup \{i\})$ and continue with the sub-partitioning to obtain a solution \tilde{x} . This solution is then compared with $\bar{x}(A, B)$. The one with the higher value is the solution for (IFD).

At first sight, Proposition 4.2 is more powerful than Proposition 4.1 since it binds more items to the players, thus making the problem simpler. Using this result, however, may result in the loss of some solutions. Part (a) of the statement does not prevent $x^*(A \cup \{i\})$ from being a possible solution for $S(A, B)$ (and a symmetrical conclusion holds for part (b)). So, if the goal is to capture all the solutions for (IFD), Proposition 4.1 is the one to choose³.

The problem remaining after the elimination test has been carried out is called the *reduced problem*. Note that the discriminating value λ is the same for the reduced problem as well as for the original problem.

4.2 The algorithm

All the elements are set to formulate a branch-and-bound algorithm for the maximin problem with indivisible items (IFD). The algorithm follows the general scheme for branch-and-bound, where the original problem $S(\emptyset, \emptyset)$ is recursively split into a series of constrained problems with some of the items assigned in advance to one player or the other. As usual for this kind of algorithms, it is convenient to represent the splitting process with a tree graph. When a subproblem cannot yield any more candidates for the solution of the original problem, the branch corresponding to that subproblem is cut (or pruned) and no other branch generates from that node of the tree.

The general framework is adapted to the peculiar features of the problem in question. For instance, the linear relaxation of each subproblem has a twofold purpose: on one hand it gives an upper bound for the value of the integer solution, but when the solution for the linear relaxation is not integer, it also suggests how to operate the splitting, by assigning the item corresponding to the unique fractional component to one player or the other.

In building the tree, several integer solutions are met and the best of them (in terms of objective function) are recorded. Here we are interested in finding all the solutions to (IFD). Therefore \bar{X} will denote the set of best solutions met so far, while \bar{z} is their common value.

Each subproblem $S(A, B)$ may have three different labels attached to it: “new”, “open” or “close”: a subproblem is new when its linear relaxation has not been computed yet; once the computation occurs, the problem is open or close depending on whether the solution for the relaxation is integer or not. Furthermore, a subproblem may also be closed when its upper bound is smaller than the best current admissible solution. Open problems are split according to the above mentioned rule. The algorithm ends when all the subproblem are closed.

The algorithm runs as follows:

³The use of Proposition 4.2 is explained with more detail in the longer version of the paper.

Initialization. Set $\bar{X} = \emptyset$ and $\bar{z} = -\infty$. Label $S(\emptyset, \emptyset)$ as new.

The generic cycle is made of the following steps

Compute bounds. For any new subproblem $S(A, B)$ perform the variable elimination test derived from Proposition 4.1 and denote with $S(A', B')$ the resulting subproblem with (possibly) more items pre-assigned to the players.

- Compute $x^+(A', B')$ and $z^+(A', B')$ using the AW-ie algorithm.
- Examine $x^+(A', B')$.
 - If $x^+(A', B')$ is integer then set $\bar{x}(A', B') = x^+(A', B')$ and $\bar{z}(A', B') = z^+(A', B')$. Label $S(A', B')$ as close.
 - If $x^+(A', B')$ has a fractional component then set $\bar{x}(A', B') = \text{rnd}(x^+(A', B'))$ with corresponding value $\bar{z}(A', B')$. Label $S(A', B')$ as close.
- Update the optimal set
 - If $\bar{z}(A', B') > \bar{z}$ then set $\bar{z} = \bar{z}(A', B')$ and $\bar{X} = \{\bar{x}(A', B')\}$.
 - If $\bar{z}(A', B') = \bar{z}$ and $\bar{x}(A', B') \notin \bar{X}$ then append this solution to \bar{X} .

List and close List the open subproblems. Close all the $S(A, B)$ such that

$$z^+(A, B) < \bar{z}. \quad (6)$$

If there is no open subproblem left, then exit the algorithm and return \bar{X} as the optimal solution set with value \bar{z} .

Choose and split Choose the open problem $S(A, B)$ with higher upper bound $z^+(A, B)$. The relaxed solution $x^+(A, B)$ has one fractional component $i \in M \setminus (A \cup B)$. Replace $S(A, B)$ (labelled close) with two subproblems $S(A \cup \{i\}, B)$ and $S(A, B \cup \{i\})$, labelling them as new. Continue with the next cycle.

Some of the rules in the algorithm may be changed. For instance another criterion may be selected to pick an open problem. A naïve motivation for the chosen rule is that the higher the bound, the more likely is the subproblem to deliver an optimal solution. Also, when $x^+(A', B')$ has a fractional component, we assign the fractional good to the player who holds more than 50% of it to get an integer admissible solution. Alternatively, we may check both options: assigning the split good to Alice and Bob, and choosing the one yielding a higher value for z . Although the latter seems a more efficient option, we prefer the simplicity of the former rule.

As noted previously, we may use a variable elimination test based on Proposition 4.2. The algorithm will be quicker, but some solutions may be left off of the solution set \bar{X} .

Since the IFD problem is NP-hard, in the worst case the algorithm may execute an exponential number of iterations to determine an optimal solution

(unless $P = NP$). To this end, let us consider the following instance of IFD: there are $m = 2n + 1$ items and both players evaluate 2 each item. In this case: the value of an optimal solution is $2n$, and corresponds to any solution which assigns n items to one player, and $n + 1$ items to the other player; the upper bound is $2n + 1$ (n variables are equal to 1, one variable is equal to 0.5, and n variables are equal to 0); a node can be closed only if either at least $n + 1$ variables are fixed equal to 1, or at least $n + 1$ variables are fixed equal to 0. Then the enumeration tree of the branch and bound, independently to possible strategies (of searching), is totally explored till level n , i.e., the number of nodes which are examined is $2^{(n+1)}$.

At the same time, the efficiency of a branch and bound technique is directly linked to the quality of (i) the method to compute bounds for each subproblem, (ii) the method to possibly split each subproblem. In our case, (i) is given by the linear programming relaxation of each subproblem, and (ii) is given by the generation of two new subproblems obtained by splitting a binary variable. The proposed algorithm adopts methods which follow the ones used in the most popular (and empirically considered efficient) branch and bound algorithms for the solution of optimization 0-1 Knapsack, which is a problem very close to IFD.

When a small number of items is at stake, the algorithm can be run by humans — in the spirit of the original AW procedure. Finding methods that keep this feature for larger bundles is the subject of our current research.

5 Acknowledgements

This is an abridged version of the paper “How to allocate hard candies fairly” which is currently being submitted for publication. The authors wish to thank two anonymous referees for the careful revision of the work and the constructive remarks. We would also like to thank Steve Brams, Erio Castagnoli and Ted Hill for useful suggestions during the completion of the work. The authors are responsible for all the remaining errors.

References

- [1] Akin, E., 1995. Vilfredo Pareto cuts the cake, *Journal of Mathematical Economics* 24, 23–44.
- [2] Bertsimas, D., and J.N. Tsitsiklis, 1997, *Introduction to linear optimization*, Athena Scientific, Belmont, Massachusetts, U.S.A.
- [3] Bezáková, I., Dani, V., 2005. Allocating indivisible goods, *ACM SIGecom Exchanges* 5, 11–18.
- [4] Brams, S.J., Edelman, P.H., Fishburn, P.C., 2001. Paradoxes of Fair Division, *Journal of Philosophy* 98, 300–314.

- [5] Brams, S.J., Edelman, P.H., Fishburn, P.C., 2003. Fair division of indivisible items, *Theory and Decision* 55, 147–180.
- [6] Brams, S.J., Fishburn, P.C., 2000. Fair division of indivisible items between two people with identical preferences: Envy-freeness, Pareto-optimality, and equity, *Social Choice and Welfare* 17, 247–267.
- [7] Brams, S.J., Jones, M.A., Klamler, C., 2006. Better ways to cut a cake, *Notices of the American Mathematical Society* 53, 1314–1321.
- [8] Brams, S.J., King, D.L., 2005. Efficient Fair Division: Help the Worst Off or Avoid Envy?, *Rationality and Society* 17, 387–421.
- [9] Brams, S.J., Taylor, A.D., 1996. *Fair Division: from Cake-cutting to Dispute Resolution*, Cambridge University Press.
- [10] Brams, S.J., Taylor, A.D., 1999. *The Win-win Solution, Guaranteeing Fair Shares to Everybody*, W.W.Norton.
- [11] Cormen, T.H., Leiserson, C.H., Rivest, R.L., Stein, C., 2001. *Introduction to Algorithms*, The MIT Press.
- [12] Demko, S., Hill, T.P., 1988. Equitable distribution of indivisible objects, *Mathematical Social Sciences* 16, 145–158.
- [13] Herreiner, D., Puppe, C., 2002. A simple procedure for finding equitable allocations of indivisible goods, *Social Choice and Welfare* 19, 415–430.
- [14] Jones, M.A., 2002. Equitable, envy-free, and efficient cake cutting for two people and its applications to divisible goods, *Mathematics Magazine* 75, 275–283.
- [15] Kuhn, H.W., 1967, On games of fair division, In: Shubik M.(Ed.), *Essays in Mathematical Economics in Honor of Oskar Morgenstern*. Princeton University Press.
- [16] Legut, J., Wilczyński, M., 1988. Optimal partitioning of a measurable space, *Proceedings of the American Mathematical Society* 104, 262–264.
- [17] Papadimitriou, C.H., Steiglitz, K., 1982. *Combinatorial Optimization: Algorithms and Complexity*, Prentice & Hall.
- [18] Wolsey, L.A., Nemhauser, G.L., 1988. *Integer and Combinatorial Optimization*, Wiley-Interscience.
- [19] Vazirani, V.V., 2001. *Approximation Algorithms*, Springer-Verlag.