

Bidding for envy-freeness: A procedural approach to n -player fair-division problems

Claus-Jochen Haake¹, Matthias G. Raith², Francis Edward Su³

¹ Institute of Mathematical Economics, University of Bielefeld, P.O. Box 100131, 33501 Bielefeld, Germany (e-mail: chaake@wiwi.uni-bielefeld.de, mraith@wiwi.uni-bielefeld.de)

² Department of Economics, University of Magdeburg, P.O. Box 4120, 39016 Magdeburg, Germany (e-mail: raith@ww.uni-magdeburg.de)

³ Department of Mathematics, Harvey Mudd College, Claremont, CA 91711, USA (e-mail: su@math.hmc.edu)

Received: 6 March 2000/Accepted: 21 May 2001

Abstract. We develop a procedure for implementing an efficient and envy-free allocation of m objects among n individuals with the possibility of monetary side-payments, assuming that players have quasi-linear utility functions. The procedure eliminates envy by compensating envious players. It is fully descriptive and says explicitly which compensations should be made, and in what order. Moreover, it is simple enough to be carried out without computer support. We formally characterize the properties of the procedure, show how it establishes envy-freeness with minimal resources, and demonstrate its application to a wide class of fair-division problems.

1 Introduction

In this paper we consider problems of *fair division*, in which a group of individuals must decide how to allocate several objects (goods or burdens) “fairly” among the group’s members, given the possibility of (monetary) side-payments. A variety of situations fit into this setting: heirs inheriting an estate, employees splitting a list of duties, developers staking claims in a new frontier, or students renting a house together. In many cases, such problems can also involve additional costs or compensations for the group as a whole.

As a basic notion of fairness, we focus on *envy-freeness*, which means that

We would like to thank Steven Brams and Marc Kilgour for stimulating discussions, and we are grateful to an anonymous referee for helpful comments. Financial support by the ‘Ministerium für Wissenschaft und Forschung, NRW’ (Raith) and the ‘Graduiertenkolleg’ Mathematical Economics, University of Bielefeld’ (Su) is gratefully acknowledged.

no individual will wish to trade shares with anyone else.¹ In proving existence of envy-free solutions, particular attention has been given to the development of *constructive* proofs that yield implementable algorithms to find the desired allocation. This aspect is emphasized by Alkan et al. (1991) and Su (1999).² Because both these approaches assume non-linear utility in money, the algorithms are not finite; instead they converge on an envy-free solution. Exact solutions can be obtained using the algorithms of Aragonés (1995) and Klijn (2000) who assume that players' preferences for money are characterized by linear utility functions. This is a strong restriction, but one that seems appropriate for the present context, nevertheless.

Our approach in this paper is not only constructive, but we consider it to be *procedural* as well. In the context of fair division, we view a procedure as featuring the following characteristics: it is *intuitive*, meaning that each step must be easy to understand; it is *plausible*, meaning that each step must be simple to argue; and it is *manageable*, meaning that each step must be straightforward to compute. We see these subjective criteria as relevant for the practical implementation of a fair-division outcome, in particular when parties in real life prefer to establish fairness by themselves, rather than trust the "magic" of a computer algorithm.

Knaster and Steinhaus (1948) proposed a fair-division method that is procedural in this sense, as it is simple to apply and largely intuitive. However, as Brams and Taylor argue as they revisit this approach, the procedure does not ensure envy-freeness for more than two players. In a more recent approach, Brams and Kilgour (2001) characterize a procedure that assigns a given number of objects to the same number of players efficiently, establishing fairness through internal "market prices". The procedure is intuitive, plausible, and simple to manage. But for more than three players it also does not guarantee envy-freeness.

In the following sections we develop a compensation procedure that establishes envy-freeness for any number of players and a possibly different number of objects. We thus extend the Knaster-Steinhaus procedure, while preserving envy-freeness for any number of players. The procedure eliminates envy in *rounds* by compensating envious players. It is fully descriptive and says explicitly which compensations should be made, and in what order. The procedure mimics and thus supports a natural mediation process with the objective of implementing an envy-free outcome. We formally characterize the properties of the procedure and illustrate how it works in practice.

We assume that players can articulate their preferences over bundles of objects through monetary bids and that their utility of money is linear. In addition we impose a *qualification* condition for each person taking part in the fair division, requiring that her valuations of all bundles sum to at least

¹ The concept of envy-freeness was first used in an economic context by Foley (1967), although under a different label.

² The algorithm of Su (1999) is even *interactive* in the sense that it sequentially gives each player the opportunity to choose her favorite alternative at evolving prices.

the total cost for the group as a whole. This is not needed for an envy-free solution, but it guarantees that a player will never have to pay more than her bid.³

We consider two alternative methods for financing the necessary monetary side-payments. In the *method of ex-ante payments*, players first pay the amount they bid for the objects they receive. Everyone thus begins with the same advantage, viz. none, and there is an aggregate amount of money left to divide among the players. Under the *method of ex-post payments*, players submit payments only after the monetary compensations needed to establish envy-freeness are determined via the procedure. A distinctive feature of both methods is that the resources for compensation are generated by the group itself. Because we find ex-ante payments more intuitive for players, we focus mainly on this method, but we do provide a comparison of the outcomes.

In Sect. 2, we formulate the first step of our procedure: an assignment of objects to players that maximizes the sum of players' utilities. We call this the *utilitarian assignment*.⁴ When the bundling of objects is restricted, the utilitarian assignment will generally cause envy among the players.

In Sect. 3, we describe non-technically the individual steps of the compensation procedure which will eliminate envy by using the surplus from players' initial payments. We demonstrate its application with a numerical example that we use throughout the paper.

The formal characterization of our compensation procedure and thus a constructive existence proof for an envy-free allocation is given in Sect. 4. Here we prove that there is always at least one player who is non-envious at the start and then show how our procedure successively eliminates the envy of players who are envious of non-envious players. In contrast to the algorithms of Aragonés (1995) and Klijn (2000), our procedure does not need to keep track of all envy relations, since it uses only *maximum* envy.⁵ A further feature distinguishing our compensation procedure from Klijn's algorithm is that it requires only minimal financial resources to establish envy-freeness. The necessary amount is automatically determined through the compensations.

In Sect. 5, we study alternative methods of dividing the surplus that generally remains after envy-freeness is established.⁶ The simplest way to obtain

³ Brams and Kilgour (2001) use a similar, but stricter constraint, basically for the same purpose.

⁴ The same starting point is chosen by Steinhaus (1948), Aragonés (1995), or Brams and Kilgour (2001). A notable exception is Klijn (2000).

⁵ Of course all envy relations are assessed by players within the procedure in order to determine maximal envy; but the non-maximal assessments are not retained for the procedure.

⁶ There will generally be an infinite number of envy-free outcomes to choose from. And, as Tadenuma and Thomson (1991) verify, there exists no proper sub-solution satisfying the notion of consistency. A sub-solution is *consistent* in their sense if it allocates the objects and money received by a subset of players in the same way as the solution assigns total resources to all players.

a unique outcome is to divide the surplus equally. However, this always leads to an outcome on the boundary of the set of envy-free prices that favors a particular player. We therefore propose an alternative method that implements a unique outcome, generally in the interior of the envy-free set, thus treating players more symmetrically.

In Sect. 6, we allow the initial assignment of exogenously given bundles to be inefficient (i.e., non-utilitarian with respect to those bundles). We find this aspect crucial, because an efficient assignment may be difficult to find in practice if the fair-division problem involves many players and objects. We show how our compensation procedure is easily adapted to non-utilitarian initial assignments. We use the fact that envy-freeness requires an efficient assignment of bundles (cf. Svensson 1983). If the assignment is inefficient, our envy-reducing procedure will create an envy cycle indicating a cyclical trade of bundles that increases the sum of players' utilities; this is analogous to the permutation procedure of Klijn (2000). Strict application of the compensation procedure thus produces cyclical trades of bundles leading to an efficient assignment, and then establishes envy-freeness as described above.

Section 7 concludes with some practical considerations. The Appendix contains numerical examples that illustrate the analysis.

2 Characterization of a utilitarian assignment

We consider a group of players $I = \{1, \dots, n\}$ who wish to assign a set of objects $K = \{1, \dots, m\}$ among themselves in an envy-free fashion.

Assumption 1. *Players value bundles of objects in a common divisible unit of account, e.g., money.*

Each player $i \in I$ can express her valuations of bundles $B_i \subseteq K$ of objects through (monetary) *bids*, which we characterize by functions $b_i : 2^K \rightarrow \mathbb{R}$. Note that Assumption 1 does not say anything specific about the relationship between a player's valuation of a bundle and her valuations of the individual objects contained within. In specific cases, however, players may have additively separable preferences over the objects in K , such that the value of a bundle simply equals the sum of values of the individual objects, i.e., $b_i(B_j) = \sum_{k \in B_j} \tilde{b}_i(k)$ ($i, j \in I$), where $\tilde{b}_i : K \rightarrow \mathbb{R}$ denotes a player's bid for specific objects. Moreover, as we will see below, when players have linear preferences over sub-divisions of objects that are divisible, the outcome of our fair-division procedure is the same whether or not K contains divisible objects.

We assume no specific relationship between the number of objects and the number of players – this can be included through additional restrictions (addl. restr.) on the objects. The group's assignment determines the *bundle* $B_i \subseteq K$ of objects that each player i receives. Possible assignments are characterized by

$$\mathcal{B} = \left\{ (B_1, \dots, B_n) \mid B_i \subseteq K, B_i \cap B_j = \emptyset, \right. \\ \left. \bigcup_{i \in I} B_i = K, (+\text{addl. restr. on } B_i) \right\}.$$

An assignment $B \in \mathcal{B}$ thus groups the m objects of K into n separate bundles B_i without dividing them; if there are fewer objects than players, B will necessarily also include empty bundles.

The set of assignments \mathcal{B} may be further restricted by specific requirements for the individual bundles. In the simplest case, the additional restriction may just be an *exogenously* given bundling of objects. For an *endogenous* bundling of objects, one could specify that all players receive the same number of objects ($|B_i| = m/n$), or that each player receives a minimum number of objects ($|B_i| \geq \underline{m}$, where $\underline{m} \leq m/n$). Or, in a different context, assume that the objects are distinct territories in a geographical region. One may then wish to have the territories in each bundle be connected in some specific form, e.g., lying within a single sub-region. Generally, the set of assignments \mathcal{B} can include any restriction on the objects of K that is player-anonymous; in particular, there are no restrictions of the form “Player i must (or must not) receive object k .” Indeed, with an exogenously imposed restriction of this type, no procedure can guarantee envy-freeness. However, if the players themselves wish to give a particular player special attention, they can express this directly via their preferences over the objects.

We view the assignment of the m objects to the n players as a *joint venture*, for which there is a total cost C (measured in the common unit of account) that must also be divided among the players. Denoting by c_i the *contribution* that is to be paid by the player receiving bundle B_i , this implies $\sum_{i \in I} c_i = C$.

Assumption 2. *Players have linear preferences over values measured in the common unit of account.*

Under Assumptions 1 and 2, we can characterize players’ preferences through quasi-linear utility functions $u_i : 2^K \times \mathbb{R} \rightarrow \mathbb{R}$, with

$$u_i(B_j, c_j) = b_i(B_j) - c_j, \quad i, j \in I.$$

In order to implement an efficient outcome, our fair-division procedure is based on an assignment that maximizes the (unweighted) sum of players’ utilities. We characterize such a *utilitarian assignment* $\bar{B} \in \mathcal{B}$ by

$$\begin{aligned} \bar{B} &\in \arg \max_{B \in \mathcal{B}} \sum_{i \in I} u_i(B_i, c_i) \\ &= \arg \max_{B \in \mathcal{B}} \sum_{i \in I} b_i(B_i) - C \\ &= \arg \max_{B \in \mathcal{B}} \sum_{i \in I} b_i(B_i). \end{aligned}$$

The utilitarian assignment \bar{B} yields the maximum sum of players’ bids,

which we denote by M . Let \bar{B}_i denote the bundle assigned to player i ; then $\sum_{i \in I} b_i(\bar{B}_i) = M$.

Regardless of the total cost C of the joint venture, the utilitarian assignment \bar{B} endogenously bundles the objects of K , while acknowledging the additional restrictions, and assigns these bundles to the individual players.⁷ If there are no additional restrictions on the bundles, the utilitarian assignment is easy to implement when preferences are additively separable: simply assign each object to the player who values it most (if there are several, choose one player arbitrarily). However, in general assignments will be complicated by the additional restrictions specifying how bundles are to be created. When bundles are given exogenously, we call an assignment *efficient* if no re-assignment of the bundles yields a larger sum of bids. The utilitarian assignment thus allocates given bundles efficiently among players.

We characterize a utilitarian allocation as *envy-free* if no player values the bundle of any other player (net of its cost) higher than her own bundle (net of its cost):

$$u_i(\bar{B}_i, c_i) \geq u_i(\bar{B}_j, c_j), \quad i, j \in I.$$

We wish to determine an envy-free pricing of utilitarian bundles, with prices that sum to the total cost C of the joint venture, such that no player pays more than she thinks her bundle is worth. The procedure described in the next section will accomplish this, if we impose the following additional requirement.

Assumption 3. *The sum of each player's bids for all the bundles of a utilitarian assignment is at least equal to the total cost, i.e., $\sum_{j=1}^n b_i(\bar{B}_j) \geq C, \forall i \in I$.*

Assumption 3 can be seen as an individual *qualification constraint* for each group member. If the objects to be distributed are assigned across several players, then player i is qualified if, by teaming with other players of identical preferences, this group of players would be able to afford the joint venture. As the procedure will show, the qualification constraint is not required to produce envy-freeness, but it guarantees that no player will pay more than her bid.

3 The compensation procedure

Our procedure with ex-ante payments begins by having each player contribute the amount that they bid for their assigned bundle, yielding M dollars from which the cost is paid. The remaining *surplus* $M - C$ will be returned to the players in the form of discounts in a way which will guarantee envy-freeness. In each round of the compensation procedure, discounts are determined on the basis of players' assessments, and then assessments are revised taking discount changes into account.

⁷ Note that, when players have linear preferences over divisible objects, the utilitarian assignment would only divide an object if the value added by its inclusion in a player's bundle is the same for two or more players. In this case we may just as well assume that the object is fully assigned to just one player.

Let a_{ij} denote Player i 's assessment of the value of Player j 's bundle minus its cost:

$$a_{ij} = b_i(\bar{B}_j) - c_j = b_i(\bar{B}_j) - b_j(\bar{B}_j) + d_j, \quad i, j \in I, \tag{1}$$

where d_j is the discount that Player j has received during the procedure (at the start $d_j = 0$). We call $A = (a_{ij})$ the *assessment matrix*. Note that if $a_{ii} < a_{ij}$ then Player i will experience envy for Player j . Without additional restrictions on bundles, the utilitarian assignment simply assigns each object to the player who values it most. The assessment matrix will thus be envy-free from the start, since $a_{ii} \geq a_{ij}$. With additional restrictions, however, this will generally not be the case, and envious players will need to be compensated. The complete compensation procedure is described as follows.

The compensation procedure for a utilitarian assignment

1. Assign bundles to players using the utilitarian assignment. Each player initially contributes her bid on her assigned bundle, yielding a pool of size M from which the cost C is paid.
2. Calculate the assessment matrix. Note that there will always be at least one player who experiences no envy (see Theorem 1). If all players are non-envious, skip to Step 5.
3. Now perform a round of compensations: use the assessment matrix to identify all players whose maximum envy is directed towards a non-envious player, and compensate these individuals from the surplus by their maximum envy difference.⁸ Then recalculate the assessment matrix (but only after all the compensations have been made in this round).
4. Perform additional compensation rounds until all envy is eliminated. (Theorem 2 shows that at most $(n - 1)$ compensation rounds will be needed.)
5. The sum of the compensations made in Steps 3 and 4 is minimal (see Theorem 3), and it will never exceed the surplus $M - C$ (see Theorem 4). Therefore distribute any remaining surplus in a way that maintains envy-freeness; e.g., one could simply divide it equally among all players. (Sect. 5 discusses an alternative method for post-envy allocation of the remaining surplus.)

To illustrate we give an example. Suppose there are four players (denoted P_i) who submit bids for a joint venture that has a total cost of $C = 100$. The utilitarian assignment determines four bundles (denoted \bar{B}_i) for which players have the valuations given in Table 1.

The bids in the utilitarian assignment (the framed entries along the diagonal of Table 1) are collected as initial payments. Since they sum to 145, after paying the cost of 100, there is a surplus of 45 left to return to the players in the form of discounts. The assessment matrix can be computed by subtracting

⁸ Alternatively, one could compensate *all* envious players. However, this would require more compensations in each round, but not fewer rounds.

Table 1. Players' bids for bundles

	\bar{B}_1	\bar{B}_2	\bar{B}_3	\bar{B}_4
P1	50	20	10	20
P2	60	40	15	10
P3	0	40	25	35
P4	50	35	10	30
Initial payment	50	40	25	30

Table 2. The initial assessment matrix

	P1	P2	P3	P4
P1	0	-20	-15	-10
P2	10	0	-10	-20
P3	-50	0	0	5
P4	0	-5	-15	0
Discounts	0	0	0	0

Table 3. The modified assessment matrix

	P1	P2	P3	P4
P1	0	-10	-10	-10
P2	10	10	-5	-20
P3	-50	10	5	5
P4	0	5	-10	0
Discounts	0	10	5	0

the diagonal entry from each column. In Table 2, row i then shows Player i 's assessment of Player j 's bundle. We keep track of discounts in a separate row.

The assessment matrix in Table 2 shows (by comparing entries in each row) that Player 2 envies Player 1 ($a_{22} < a_{21}$) and Player 3 envies Player 4 ($a_{33} < a_{34}$). Therefore we must compensate Player 2 by giving her a discount of 10, and Player 3 a discount of 5. To recalculate the assessment matrix, we may add 10 to column 2 and add 5 to column 3. The new assessment matrix is given in Table 3.

Now both Player 3 and Player 4 envy Player 2, and Player 2 feels tied with Player 1 (who remains non-envious). We must compensate Player 3 and Player 4 by giving them both additional discounts of 5. Adding 5 to both columns 3 and 4, we obtain Table 4.

Table 4. The envy-free assessment matrix

	P1	P2	P3	P4
P1	0	-10	-5	-5
P2	10	10	0	-15
P3	-50	10	10	10
P4	0	5	-5	5
Discounts	0	10	10	5

Now all envy has been eliminated, since each diagonal element is the largest entry in its row. The discounts used 25 units of the surplus, and the remaining surplus of 20 can be equally divided among the four players, yielding total discounts given by $d = (5, 15, 15, 10)$. This gives final envy-free costs of $c = (45, 25, 10, 20)$.

It is important to note that, in formulating the compensation procedure, we do not make any assumptions concerning the signs of players' bids b_i or their contributions c_i . Therefore, our procedure can also be applied to situations where the objects are *burdens* for which the group as a whole receives a *compensation* ($C < 0$). Players' negative bids are then requested payments that express their disutility of accepting these burdens, and Assumption 3 states that a player is qualified if her demands for bearing all burdens do not exceed the total compensation, i.e., $-b_i(K) \leq -C$. We provide an example for the division of burdens in the Appendix.

More generally, our procedure can be applied to fair-division problems that involve both goods *and* burdens. For example, a group of individuals that decides to share a house cannot only derive a fair allocation of rooms and rents, but they can also include all the (group's) chores that come with the house (e.g., lawn mowing, cleaning, cooking, etc.). Of course, if there is no extra compensation, the qualification constraint (Assumption 3) becomes more binding as chores are added. But this is only plausible – in order to qualify, a housemate must not only be willing to pay the necessary rent, she must also be willing to perform the necessary chores.

4 Properties of the compensation procedure

We now show that our procedure works as indicated. One property of the assessment matrix A is crucial in all that follows. We define a *permutation sum* of an $n \times n$ matrix A to be any sum of the form $\sum_i a_{i\pi(i)}$, where $\pi : I \rightarrow I$ is a permutation of n elements. Thus a permutation sum picks one element of each column and row and forms their sum.

Lemma 1. *At any step of the procedure, the largest permutation sum of A occurs along the diagonal.*

Proof. We check:

$$\begin{aligned} \sum_i a_{i\pi(i)} &= \sum_i (b_i(\bar{B}_{\pi(i)}) - b_{\pi(i)}(\bar{B}_{\pi(i)}) + d_{\pi(i)}) \\ &= \sum_i b_i(\bar{B}_{\pi(i)}) - M + \sum_i d_i \leq \sum_i d_i, \end{aligned}$$

which follows from the definition of M , the maximum sum of bids. The inequality is clearly an equality when π is the diagonal assignment. \square

We keep track of “chains” of envy with the following notation. We write $i \rightarrow j$ if j is the player that i envies the most (if there are several, pick one arbitrarily). In this case $a_{ii} < a_{ij}$ and a_{ij} is the largest entry in row i . The single arrows thus only indicate maximum envy relations. We use a double arrow $i \Rightarrow j$ if i envies no one but feels tied with j and this tie was the result of an earlier compensation. Thus $a_{ii} = a_{ij}$, and these are the largest entries in row i . Note that double arrows only keep track of *created* ties, and are *not* used for ties occurring “by coincidence” (e.g., where $a_{ii} = a_{ij}$ but this was not the result of a previous compensation). In the sequel when we refer to “arrow” we shall mean either a single or double arrow unless explicitly specified.

We form a directed graph G in which the vertices represent players and the edges are given by the arrow relations between players.⁹ (We shall speak of “players” and “vertices” interchangeably in all that follows.) Throughout the procedure, the directed graph G will change. Keeping track of how G evolves is the key to showing that the procedure terminates.

Lemma 2. *At any step in the procedure, the directed graph G contains no cycles.*

Proof. If there were a cycle of single arrows, say $i_1 \rightarrow i_2 \rightarrow \dots \rightarrow i_k \rightarrow i_1$, then

$$\begin{aligned} a_{i_1 i_1} &< a_{i_1 i_2} \\ a_{i_2 i_2} &< a_{i_2 i_3} \\ &\vdots \\ a_{i_k i_k} &< a_{i_k i_1}, \end{aligned} \tag{2}$$

and by adding these relations one would find that

$$a_{i_1 i_1} + a_{i_2 i_2} + \dots + a_{i_k i_k} < a_{i_1 i_2} + a_{i_2 i_3} + \dots + a_{i_k i_1}, \tag{3}$$

which augmented by the other diagonal terms would contradict Lemma 1.

A nearly identical argument can be used for cycles in which some (but not all) of the arrows are double; some of the envy inequalities in (2) would become equalities, but the inequality (3) would remain strict as long as there were a single arrow in the cycle.

The only other possibility is a cycle consisting entirely of double arrows. But this cannot arise, because double arrows only originate from single arrows

⁹ It is important to note that our directed graph G is simpler than the graphs constructed by Aragonés (1995) and Klijn (2000), since our procedure only keeps track of maximum envy.

(as compensations are made to envious players). Thus if there were a cycle of double arrows, by stepping backwards through the procedure one would find a prior round in which that cycle contained a single arrow – a possibility that was ruled out above. \square

The *outdegree* of a vertex represents the number of arrows that originate at a vertex. Because of the way we defined arrows, every vertex in G has an outdegree of at most 1. Hence there is a uniquely defined *path* that “flows” from any given vertex. Since G cannot have any cycles at any step in the procedure, we deduce that G must always be a disjoint set of directed trees, each of which has a unique *root*, a vertex of outdegree 0 to which all other vertices flow. Roots of trees correspond to non-envious players who have not yet experienced envy.

Theorem 1. *At the start there is at least one player who will remain non-envious throughout the entire procedure.*

Proof. A vertex with no arrow is a root, and a vertex with a single or double arrow corresponds to a player who experiences or has experienced envy. It follows that the outdegree of a vertex can never decrease – once a root earns an arrow it can never become a root again.

If by the end of the procedure there were no roots, then any path following the arrows in G would eventually cycle, contradicting Lemma 2. Thus there is a root which must have been a root throughout the entire procedure, corresponding to a person who remains non-envious. \square

Because every vertex in a tree has a unique path to its root, we may classify vertices in G by “levels” – a level 0 vertex is a root, and a level k vertex is one that is k arrows away from the root of its tree. A vertex i is said to be an *ancestor* of vertex j in the tree if there is a chain of arrows flowing from j to i . Thus the root of a tree is an ancestor of every other vertex in the tree. The *ancestral path* of vertex i is the set of all vertices on the path from i to its root – including the root, but not including i . (Note that the ancestral path does not depend on the type of arrows along the path.) Vertices may change levels as G evolves throughout the procedure.

Lemma 3. *During a compensation round, if a player P changes her ancestral path, the new path must contain a newly compensated player Q between P and the root. Thus in the new graph, P will have a higher level number than Q .*

Proof. During a compensation round, the only new envy that can be introduced is directed at players receiving compensations during that round. Thus if player P or any of her ancestors changes her maximum envy (hence the direction of her arrow) to a newly compensated player Q , then after the compensation round, the path flowing from P must pass through Q . From the definition of level this means that P will have a higher level number than Q . \square

Lemma 4. *After k compensation rounds, there are no envious players on levels 0 through k .*

Proof. We prove this by induction on k . For $k = 0$ (before any compensations) the statement trivially holds, since there are by definition no envious roots.

Now assume that the lemma holds for k . We show that it also holds for $k + 1$. By the inductive hypothesis, after k compensation rounds all envious players must be on levels $(k + 1)$ or higher. A subset of these will be compensated during the $(k + 1)$ -st round; in particular this must include all the envious players on level $(k + 1)$, since their maximum envy is for non-envious players on level k .

So consider a player P anywhere in G . If P 's ancestral path did not change as a result of the $(k + 1)$ -st compensation round, then she remains on the same level. Moreover, if she is on level $(k + 1)$, then she must also now not be envious (because she was either compensated or was not envious to begin with).

If P 's ancestral path did change as a result of the $(k + 1)$ -st compensation round, then we show her new level number must be $(k + 2)$ or greater. Consider the new path that flows from P ; by Lemma 3, it contains some newly compensated player. Thus it makes sense to speak of the newly compensated player on P 's path who is *closest* to the root; call this player Q . Q 's ancestral path cannot have been changed by the last compensation round (otherwise Lemma 3 would have produced some other newly compensated player closer to the root than Q). So Q 's level was unchanged, and being a newly compensated player, Q must have had level number $(k + 1)$ or greater. Since P has a higher level number than Q , P must have a level number of $(k + 2)$ or greater in the new graph.

Thus if any player changed levels as a result of the $(k + 1)$ -st compensation round, they must now be at level number $(k + 2)$ or greater. The players who remain on levels 0 through k were non-envious before the round and must still be non-envious, while the ones remaining on level $(k + 1)$ have been compensated and are also non-envious. \square

Theorem 2. *The procedure requires no more than $n - 1$ compensation rounds to eliminate envy.*

Proof. Because G has only n vertices, and every tree in G must have a root, no vertex can be on level n or greater. Hence, by Lemma 4, all players will be non-envious after at most $(n - 1)$ compensation rounds. \square

It is not yet clear what effect the initial assignment has on the resulting envy-free discounts. Moreover, when the utilitarian assignment for n given bundles is not unique, players have multiple choices for the starting point of the compensation procedure. Remarkably, the outcome is not affected by the choice of assignment. We first prove two lemmas needed to establish this result.

Lemma 5. *If there are two utilitarian assignments involving the same bundles, then a vector of discounts which yields envy-free assessments under one assignment will also be envy-free under the other assignment.*

Proof. We have

$$a_{ij} := b_i(\bar{B}_j) - b_j(\bar{B}_j) + d_j, \quad i, j \in I.$$

Let $\pi : I \rightarrow I$ be the permutation that transforms the original utilitarian assignment into the other one; i.e., in the second assignment Player i receives the bundle that Player $\pi(i)$ would receive in the original assignment. Envy-freeness of the discounts d_i in particular yields $a_{ii} \geq a_{i\pi(i)}$.

Taking the sum over differences yields

$$\begin{aligned} \sum_i (a_{ii} - a_{i\pi(i)}) &= \sum_i (d_i - d_{\pi(i)}) + \sum_i (b_{\pi(i)}(\bar{B}_{\pi(i)}) - b_i(\bar{B}_{\pi(i)})) \\ &= 0 + M - M = 0. \end{aligned}$$

Since each addend on the left-hand side is non-negative, and their sum is zero,

$$a_{ii} = a_{i\pi(i)}, \quad \forall i \in I. \quad \square$$

Said another way, Lemma 5 shows that the compact convex set of all possible envy-free discount vectors is independent of the utilitarian assignment.¹⁰ Note, however, that Lemma 5 says nothing specific about the discount vector induced by the compensation procedure.

Lemma 6. *The compensation procedure yields a unique minimal vector of non-negative discounts that make a given utilitarian assignment envy-free.*

Proof. Given a utilitarian assignment, the compensation procedure yields a vector of player discounts (d_i) which makes every player envy-free. If there were some other vector of (non-negative) envy-free player discounts (m_i) which was smaller for some player, then we obtain a contradiction.

Suppose $m_{k_0} < d_{k_0}$, for player k_0 . Thus d_{k_0} is strictly positive (because $m_{k_0} \geq 0$), i.e., player k_0 was compensated sometime during the compensation procedure. This implies that in the final envy graph at the conclusion of the compensation procedure, Player k_0 is not at the root of her tree. Let $k_0, k_1, k_2, \dots, k_l = r$ be the path from k_0 to the root r of her tree. Since each of the arrows in the final envy-graph are double arrows, $a_{k_i k_i} = a_{k_i k_{i+1}}$ for all $0 \leq i \leq l - 1$. Using (1) to express this in terms of the bids and discounts, we have

$$d_{k_i} = b_{k_i}(\bar{B}_{k_{i+1}}) - b_{k_{i+1}}(\bar{B}_{k_{i+1}}) + d_{k_{i+1}},$$

for all $0 \leq i \leq l - 1$. But the (m_i) , being an envy-free vector of discounts, must satisfy

$$m_{k_i} \geq b_{k_i}(\bar{B}_{k_{i+1}}) - b_{k_{i+1}}(\bar{B}_{k_{i+1}}) + m_{k_{i+1}},$$

since the left side is what Player k_i would receive under these discounts and the right side is what Player k_i believes Player k_{i+1} would receive. Subtracting these two equations we obtain

¹⁰ Lemma 5 corresponds to Aragonés' (1995) Lemma 4.

$$d_{k_i} - d_{k_{i+1}} \leq m_{k_i} - m_{k_{i+1}}.$$

By summing this over all $0 \leq i \leq l - 1$, we have

$$d_{k_0} - d_{k_l} \leq m_{k_0} - m_{k_l}.$$

But $d_{k_l} = 0$ because k_l was the root, i.e., a player not compensated throughout the entire procedure. And since the discount m_{k_l} is non-negative, we have

$$d_{k_0} = d_{k_0} - d_{k_l} \leq m_{k_0} - m_{k_l} \leq m_{k_0}.$$

This contradicts the fact that $m_{k_0} < d_{k_0}$. □

Together Lemmas 5 and 6 allow us to establish the following practical result.

Theorem 3. *For given bundles, the outcome of the compensation procedure yields a unique minimal vector of player discounts which does not depend on the utilitarian assignment chosen.*

Proof. Suppose there were two utilitarian assignments – the non-prime and the prime assignment. The compensation procedure applied to each assignment yields corresponding vectors of player discounts (d_i) and (d'_i) .

Note that the discounts (d'_i) are also envy-free discounts for the non-prime assignment (Lemma 5), and because the (d_i) are minimal for the non-prime assignment (Lemma 6), we must have $d_i \leq d'_i$ for all i . Similarly, the (d_i) are also envy-free discounts for the prime assignment, and minimality of the (d'_i) for the prime assignment implies that $d'_i \leq d_i$ for all i . Therefore $d_i = d'_i$ for all i . □

The sum of discounts thus yields the minimal amount of money required for envy-freeness. With n bundles to allocate among the individual players, multiple utilitarian assignments do not cause a coordination problem, because they all lead to the same discounts. Hence there is no problem of choosing an appropriate starting point for the compensation procedure as long as the assignment maximizes the sum of players' bids.

We have shown that this procedure will terminate with envy-free costs without having used the individual qualification condition (Assumption 3). The only need for this condition is to ensure that the surplus is never exceeded at the end of the compensation rounds. (If the surplus were exceeded at the end, one could still obtain envy-free prices by charging all players equally for the overdraft, but then some players might end up paying more than what they bid on their bundle.)

Theorem 4. *If each person meets the qualification condition (Assumption 3), then by the end of the procedure, the compensations will not have exceeded the surplus.*

Proof. Since the prices obtained at the end of the procedure are envy-free, $a_{ii} \geq a_{ij}$ for all i, j . Using (1) this implies

$$d_j - d_i \leq b_j(\bar{B}_j) - b_i(\bar{B}_j), \tag{4}$$

for all i, j . Summing the above equation over all j , one obtains that for all i ,

$$\sum_j (d_j - d_i) \leq M - \sum_j b_i(\bar{B}_j). \tag{5}$$

Choose any player i who was not compensated throughout the entire procedure (there must be at least one, by Theorem 1). Since $d_i = 0$, and since the sum of the bids of any player is C or greater, we must have that

$$\sum_j d_j \leq M - C,$$

which shows that the sum of the compensations does not exceed the surplus. □

5 Dividing the remaining surplus

According to Theorem 4, after the compensation procedure has established envy-freeness, the remaining surplus S is given by

$$S = M - C - \sum_{j \in I} d_j \geq 0.$$

With $S > 0$, there is surplus left to distribute among the players. The distribution schemes that are of interest here are those that maintain envy-freeness. There is a convex and compact set of envy-free discounts to choose from.¹¹ We consider two alternative methods for implementing a unique solution.

Equal distribution of the surplus or “ex-post equal payments”

With an envy-free assessment matrix at the end of the compensation procedure, no envy will be created if all entries in the matrix are increased by the same amount. This is easily achieved through an equal distribution of the remaining surplus among all players. Denoting players’ final discounts under this equal distribution scheme by d_i^e , this gives

$$d_i^e = d_i + \frac{1}{n}S. \tag{6}$$

The equal distribution of the remaining surplus was demonstrated in our example in Sect. 3. Despite the simplicity of this distribution scheme, some parties may dislike the procedural asymmetry because they pay different amounts for their bundles but receive identical shares of the remaining surplus.

Therefore, consider the following modification of the compensation procedure: parties are assigned bundles, but they *do not* pay in advance. Instead, a (hypothetical) mediator finances the compensation procedure and charges the group afterwards for total compensations and the cost C . The mediator

¹¹ The set of envy-free prices is given by a convex polyhedron characterized by n^2 inequalities that can be derived from the envy-free assessment matrix.

lets each party pay an equal share of the total costs. We call this the compensation procedure with *ex-post equal payments*.

With this modification, Player i 's assessment of Player j 's bundle becomes $a_{ij} = b_i(\bar{B}_j) + \tilde{d}_j$, where \tilde{d}_j denotes the compensation under the modified procedure. Thus, in this case the initial assessment matrix is the bid matrix, so it follows that Theorems 1–3 also hold for the method of ex-post equal payments. And, similar to the proof of Theorem 4, it can be shown using the qualification condition that $\frac{1}{n}(\sum_j \tilde{d}_j + C) \leq b_i(\bar{B}_i) + \tilde{d}_i, \forall i \in I$. Hence no player pays more than what she thinks her share is worth.

Generally, the directed graphs for ex-ante and ex-post payments will feature different envy relations. Nevertheless, the seemingly different procedures exhibit an interesting equivalence: if the final directed graphs turn out to be the same and exactly one and the same player remains non-envious in both procedures, then the ex-post equal payments procedure and the ex-ante payments procedure with an equal distribution of the remaining surplus will yield the same outcome.¹²

Theorem 5. *Consider the two procedures: (1) ex-post equal payments and (2) ex-ante payments with equal distribution of the remaining surplus. Suppose that the final envy graphs of both procedures coincide and have exactly one root. Then the outcome of both procedures will be the same.*

Proof. With ex-post payments, each player $i \in I$ begins with an initial value $b_i(\bar{B}_i)$. Player i envies Player j if Player i thinks that Player j receives more than Player i . In the (final) directed graph, which we denote by \tilde{G} , $i \Rightarrow j$ then implies

$$\tilde{d}_i = b_i(\bar{B}_j) - b_i(\bar{B}_i) + \tilde{d}_j, \tag{7}$$

where \tilde{d}_i denotes the compensation under the modified procedure.

In contrast, the compensation procedure with ex-ante payments begins by having each player $i \in I$ pay $b_i(\bar{B}_i)$ for her assigned bundle. Player i envies Player j if Player j pays less for her bundle than what Player i thinks it is worth. In the (final) directed graph G , $i \Rightarrow j$ implies

$$d_i = b_i(\bar{B}_j) - b_j(\bar{B}_j) + d_j. \tag{8}$$

Coincidence of the final graphs enables us to subtract (8) from (7) and to obtain:

$$b_i(\bar{B}_i) + \tilde{d}_i - d_i = b_j(\bar{B}_j) + \tilde{d}_j - d_j.$$

Assume that Player r is a root of the directed graphs G and \tilde{G} . Since Player r does not receive a compensation under either procedure, $d_r = \tilde{d}_r = 0$, hence

$$b_i(\bar{B}_i) + \tilde{d}_i - d_i = b_r(\bar{B}_r) \tag{9}$$

¹² This is the case for our example given in Table 1. The ex-post payments procedure yields final discounts $\tilde{d} = (0, 20, 35, 25)$. With $C = 100$, there is a total of 180 units to be shared equally by all four players. The final cost for each player is thus the same as under the ex-ante method.

for every player i in the directed graph leading (with double arrows) to Player r . Summing up both sides in (9) we obtain in particular

$$\frac{1}{n} \sum_j [b_j(\bar{B}_j) + \tilde{d}_j - d_j] = b_r(\bar{B}_r). \tag{10}$$

Note that the final value that Player i receives under the compensation procedure with *ex-ante payments and an equal distribution of the remaining surplus* is

$$\begin{aligned} & b_i(\bar{B}_i) - b_i(\bar{B}_i) + d_i + \frac{1}{n} \left[\sum_j b_j(\bar{B}_j) - C - \sum_j d_j \right] \\ &= b_i(\bar{B}_i) - [b_i(\bar{B}_i) - d_i] + \frac{1}{n} \sum_j [b_j(\bar{B}_j) - d_j] - \frac{1}{n} C. \end{aligned}$$

This is the value of Player i 's bundle, minus its (discounted) cost, plus an equal share of the total payments, minus an equal share of the total costs.

Consider now the final value that Player i receives under the compensation procedure with *ex-post payments and an equal distribution of total compensations and total costs* C :

$$b_i(\bar{B}_i) + \tilde{d}_i - \frac{1}{n} \sum_j \tilde{d}_j - \frac{1}{n} C.$$

The final value is the value of the bundle, plus additional compensation, minus an equal share of all costs.

By comparing their final values, one can see that both procedures lead to the same outcome if and only if

$$b_i(\bar{B}_i) + \tilde{d}_i - d_i = \frac{1}{n} \sum_j [b_j(\bar{B}_j) + \tilde{d}_j - d_j].$$

Equations (9) and (10) show that this condition is satisfied if the final envy-graphs of both procedures are the same and possess a single root. \square

But even when there is a difference in the outcomes (due to different envy relations), the compensation procedure with ex-post equal payments establishes envy-freeness with minimal resources and thus minimal side-payments between players. Afterwards each player is charged an equal amount, just enough to cover the *cost of envy-freeness* plus the *cost of the joint venture*, so that there is no remaining surplus to be distributed. Hence, if the sole objective is to implement a unique envy-free outcome, the method with ex-post payments has a practical advantage. We demonstrate the application of this procedure in Sect. 6.

The average discount method or “The average biased mediator”

An equal distribution of the remaining surplus may not be the most plausible approach if the set of envy-free discounts is asymmetric in the sense that

players' maximum possible discounts under envy-freeness differ. The asymmetry becomes apparent if the remaining surplus is used to maximize the discount of a specific player i , while raising the discounts of the other players just enough to maintain envy-freeness. If this is done with each of the n players, one obtains n extreme discount distributions, each favoring a specific player. The *average discount method* takes the average of these n extreme surplus distributions, yielding a unique outcome, generally in the interior of the envy-free set.¹³

Focusing on a specific player $i \in I$, we denote by d_{ji} Player i 's own maximum discount and the minimum corresponding discounts of all other players $j \in I$ by d_{ji} . At the start, all players' discounts are as given by the compensation procedure, i.e., $d_{ji} := d_j$, $j \in I$.¹⁴ The *extreme* discounts d_{ji} are then updated according to the following four-step algorithm.

The average discount method

- (i) Begin by placing Player i in the set D of players whose discounts are to be increased. Formally: $D := \{i\}$.
- (ii) Add to the set D all the players who are not in D but feel tied with some player in D . Note that these are the players who would become envious if the discounts of the players in D were increased. Continue looking for additional players to be included in D until only those players are left (outside of D) who value the discounted bundle of every player in D less than their own. Formally: $D^+ := D \cup \{h \in I \setminus D \mid a_{hj} = a_{hh}, j \in D\}$. If $D^+ \neq D$, then $D := D^+$, and repeat Step (ii).
- (iii) If the set D contains all the players of I , then distribute the remaining surplus equally among the players. Otherwise, determine how far the discounts of the players in D can be raised without creating envy for any player who is not in D . The discounts of the players in D can then be increased up to this maximum amount as long as there is enough surplus left. Update the discounts of the players in D , recalculate the corresponding columns of the assessment matrix, and recalculate the remaining surplus. Formally: if $D = I$, then increase the discounts of all players by $d^+ := \frac{S}{|D|}$; else determine $d^+ := \min\left(\min_{h \in I \setminus D} \{a_{hh} - a_{hj}, j \in D\}, \frac{S}{|D|}\right)$. Update: $d_{ji} := d_{ji} + d^+$; $a_{ij} := a_{ij} + d^+$, $\forall j \in D$; $S := M - C - \sum_{j \in I} d_{ji}$.
- (iv) If there is no surplus left (which happens if $d^+ = S/|D|$), quit. Otherwise, if some surplus remains, we need to update D by returning to Step (ii). (The updated discounts from the last step will have created players, not in D , who feel tied with players belonging to D .) Formally: if $S = 0$, quit. If $S > 0$, return to Step (ii).

¹³ More generally, one could also maximize the discounts of every subset of players, thus tracing out *all* the corners of the polyhedron of envy-free prices, and then take the midpoint of all outcomes.

¹⁴ We use “:=” as an assignment operator to update variables in the following way: the value of the term on the right-hand side is assigned to the variable on the left-hand side.

After the algorithm has been applied for each player $i \in I$, determine the average extreme discount of Player i :

$$d_i^a = \frac{1}{n} \sum_{j \in I} d_{i|j}. \tag{11}$$

We demonstrate the calculation of average discounts by using our example of Sect. 3. The final outcome of the compensation procedure with ex-ante payments is shown in Table 4, with a remaining surplus of $S = 20$ to be distributed among the four players. Beginning with Player 1, the final discounts are (5, 15, 15, 10); beginning with Player 2, the final discounts are (1.25, 16.25, 16.25, 11.25); beginning with Player 3, the final discounts are (3.75, 13.75, 18.75, 8.75); and beginning with Player 4, the final discounts are (2.5, 12.5, 17.5, 12.5). A detailed derivation of these discounts is given in the Appendix. Taking all four extreme discount distributions into account, the average discounts, given by Eq. (11), are $d^a = (3.125, 14.375, 16.875, 10.625)$, yielding final costs of $c = (46.875, 25.625, 8.125, 19.375)$.

In all four cases of individual discount maximization, as soon as Player 1 is added to D , all other players are included as well, and the remaining surplus is divided equally among the whole group. This is because Player 1 is the root of a unique tree of double arrows at the end of the compensation procedure. So if her discount is increased, all other players must benefit equally in order to avoid new envy.

More generally, when there are several roots, as soon as the last uncompensated player i with $d_i = 0$ joins D , the whole group must belong to D , i.e., $D = I$. At this point, however, the aggregate discounts cannot yet have exceeded the surplus if all players met the qualification constraint (Assumption 3) – this follows from the proof of Theorem 4.

The maximization of a specific player i 's discount according to the method above could also be interpreted as the outcome of a modified compensation procedure implemented by a *biased mediator* who favors Player i . This only requires modifying Step 3 of the compensation procedure: whenever the favored player is compensated for her maximum envy and then, if no one envies her, increase her discount *as much as possible* without raising the envy of any other player.¹⁵ Compensate each other player the same as before. Distribute the remaining surplus according to the average discount method, favoring the same player. This biased procedure can be applied for each player. The average over all biased compensations thus yields the outcome of the *average biased mediator*.

Recall that under an equal distribution of the remaining surplus, given by Eq. (6), each player received an increment of 5 to her discount. In the pre-

¹⁵ In the directed graphs used in the proofs of Lemmas 2–4, a double arrow $i \Rightarrow j$ then means that i envies no one, but feels a (weak) *advantage* over j , and this *lead* was the result of an earlier compensation. Thus $a_{ii} \geq a_{ij}$, and these are the largest entries in row i . As before, double arrows only keep track of *created* leads. Note that the biased compensations during the compensation procedure do not change the envy graph.

ceding analysis of extreme envy-free prices, we found that 5 is the maximum discount that Player 1 (the root player) can receive. In contrast, for Player 3, an increase of 5 is the minimum increment that preserves envy-freeness. The equal distribution of the remaining surplus thus implements the most favorable outcome for the initially non-vious (and therefore uncompensated) player on the boundary of the set of envy-free prices. This may be difficult to justify in practice.

By contrast, the average discount method acknowledges an additional notion of fairness: Player 1, who did not experience envy throughout the entire compensation procedure and who is always at the verge of being envied by Player 2, receives only a small share of the remaining surplus. Player 3, on the other hand, who quickly becomes envious when other players are compensated, but who is relatively far from being envied by anyone else, receives a larger share. From a practical viewpoint, the average discounts enhance the stability of the outcome: by choosing discounts in the *interior* of the set of envy-free prices, each player strictly prefers her own bundle to any other (unless, of course, some players have identical preferences).

6 Cycling to efficiency

In problems of fair division that involve only a few parties and a few objects, a utilitarian assignment will usually be easy to identify in practice. When many players are involved, the complexity of this initial step quickly rises. Clearly, if players must rely on computational assistance to perform the necessary calculations before compensations can be made, this will diminish the attractiveness of a procedure that is supposed to work without computer support.

In our previous analysis, we used a utilitarian assignment as the starting point for our procedure, thus ensuring that the envy-free allocation is also efficient. However, efficiency is not just a further desirable property of the outcome. Indeed, envy-freeness can only be achieved if the assignment of n given bundles is efficient.

Lemma 7. *Let d be a vector of envy-free discounts for a given assignment $B \in \mathcal{B}$. Then B must be an efficient assignment of the n associated bundles.*

Proof. Consider some assignment $B \in \mathcal{B}$ and an assignment $B_\pi \in \mathcal{B}$, which is obtained through a permutation π of the n bundles of B . Let d be the vector of envy-free discounts associated with the assignment B . Envy-freeness then implies $a_{ii} \geq a_{i\pi(i)}$, or, using definition (1),

$$d_i \geq b_i(B_{\pi(i)}) - b_{\pi(i)}(B_{\pi(i)}) + d_{\pi(i)}. \quad (12)$$

Since $\sum_i d_i = \sum_i d_{\pi(i)}$ and $\sum_i b_{\pi(i)}(B_{\pi(i)}) = \sum_i b_i(B_i)$, we can sum over both sides of (12) to obtain

$$\sum_i b_i(B_i) \geq \sum_i b_i(B_{\pi(i)}),$$

for any permutation π . Hence B is an efficient assignment of the n bundles. \square

So, what if the initial assignment $B \in \mathcal{B}$ is not efficient? With exogenously given bundles, the utilitarian assignment \bar{B} can be obtained from B through a permutation of bundles. This is equivalent to the setting studied by Klijn (2000) who assumes that the number of objects equals the number of players, with the additional restriction that each player must receive one object. If this is the case, Lemma 7 implies that envy-freeness can only be established for a utilitarian assignment.¹⁶

Consider an arbitrary assignment $B \in \mathcal{B}$. If B results in an efficient allocation of bundles, the compensation procedure will lead to envy-free discounts. If the assignment is not efficient, then Lemma 7 implies that the compensation procedure cannot establish envy-freeness, so it must lead to an envy cycle, i.e., during the compensation procedure the envy graph exhibits a cycle with at least one single arrow.

Theorem 6. *Let $B \in \mathcal{B}$ be an inefficient assignment of n given bundles. Applying the compensation procedure will then create an envy cycle (in the directed graph G) after at most $n - 1$ compensation rounds. Cycling bundles in the opposite direction of the arrows (re-assigning to each player in the cycle the bundle of the player she envies or previously envied) will increase the sum of players' utilities.*

Proof. We begin by proving the first claim. If the directed graph G of the inefficient assignment B contains cycles in single arrows (i.e., before compensations are made), the claim trivially holds. Assume therefore that G contains no envy cycles in single arrows. In that case there is at least one player who is non-envious.

The compensation procedure eliminates envy by compensating envious players. Throughout the procedure the directed graph G evolves, with single arrows being converted into double arrows and new single arrows emerging through the creation of new envy. In Lemmas 3 and 4 the notion of a *level* only applies to a non-cyclical graph, but otherwise both lemmas are valid regardless of whether or not the assignment is efficient. Assume by way of contradiction that an envy cycle is never created. Then the procedure must terminate in at most $n - 1$ rounds in an envy-free solution. This contradicts Lemma 7. Thus after at most $n - 1$ rounds the procedure must create a cycle (consisting of at least one single arrow and the rest double arrows) in the directed graph G .

In order to establish the second claim, consider the situation where a cycle is created, in which a non-envious player k_0 becomes envious of a newly compensated player k_l . (Note that k_0 does not have to be a root in G .) The cycle in G has the following structure:

$$k_l \Rightarrow \dots \Rightarrow k_1 \Rightarrow k_0 \rightarrow k_l, \quad 1 \leq l \leq n - 1,$$

¹⁶ This is equivalent to a result established by Svensson (1983).

where the cycle may contain more than one single arrow. Due to the compensations of (not necessarily all) ancestors of k_l in the graph G ,

$$d_{k_l} = b_{k_l}(\bar{B}_{k_l}) - b_{k_{l-1}}(\bar{B}_{k_{l-1}}) + d_{k_{l-1}}$$

$$\leq \sum_{h=1}^l [b_{k_h}(\bar{B}_{k_{h-1}}) - b_{k_{h-1}}(\bar{B}_{k_{h-1}})] + d_{k_0}.$$

However, since k_0 is now envious of k_l ,

$$d_{k_0} < b_{k_0}(\bar{B}_{k_l}) - b_{k_l}(\bar{B}_{k_l}) + d_{k_l},$$

and therefore,

$$\sum_{h=0}^l b_{k_h}(\bar{B}_{k_h}) < b_{k_0}(\bar{B}_{k_l}) + \sum_{h=1}^l b_{k_h}(\bar{B}_{k_{h-1}}),$$

which shows that a reassignment of bundles in the opposite direction of the arrows strictly increases the sum of players' utilities. □

Applying the compensation procedure to any assignment $B \in \mathcal{B}$, for which the directed graph G contains no envy cycles, will either lead to envy-freeness or to a cycle in G that allows a utility-increasing reassignment of bundles. Repeated application of the compensation procedure (restarting with zero discounts after each permutation of bundles), leads to an envy-free efficient outcome with respect to the given bundles, since the number of assignments is finite. Note that, unlike the permutation procedure of Klijn (2000), our procedure does not need to keep track of all envy relations. This is because the single arrows in the directed graph identify only maximum envy, i.e., possible cycles in non-maximum envy are simply ignored.

Since the computation of assessment matrices requires additional calculations when the assignment changes during the procedure, we use the compensation procedure with ex-post equal payments, where initial assessments are made directly on the basis of players' bids. From the previous section, we know that this procedure also satisfies Theorems 1–3. The modified compensation procedure consists of the following steps.

The compensation procedure for an arbitrary assignment

1. Consider all players to be non-envious at the start. For n exogenously given bundles, assign to each player an initial bundle. (For practical reasons, one could let each player state her “best guess” of a utilitarian assignment and then choose the assignment with the highest sum of bids as the initial starting point.) Use the bid matrix as the initial assessment matrix.
2. Use the assessment matrix to identify each player's maximum envy. If no player is envious, skip to Step 5.
3. If a previously non-envious player has become envious, check for an envy cycle. If a cycle exists, re-assign bundles by giving each player in the cycle the bundle of the player she envies (single arrow) or previously envied

- (double arrow). Have all players return any compensations they have received, and repeat Step 2 to restart the compensation procedure.
4. If there are envious players, perform a round of compensations: identify those players whose maximum envy is directed towards a non-envious player. Compensate them by their maximum envy difference, thus making them non-envious. Then recalculate the assessment matrix (but only after all the compensations have been made in this round), and repeat Step 2.
 5. Have each player pay an equal share of the total compensations and the total cost C .

We provide an example of the modified compensation procedure in the Appendix.

It is important to note that the procedure as formulated above cycles to a utilitarian assignment and establishes envy-freeness only for n exogenously given bundles. In particular, when each bundle contains only a single object, the final assignment will be utilitarian. In our more general framework, though, an assignment does not require the same number of objects as players. Consequently, an assignment is non-utilitarian not only when the n bundles of the assignment are *allocated* inefficiently, but also if the m objects are *bundled* inefficiently. Since cyclical trades under the compensation procedure involve complete bundles of objects, players may thus still wish to trade individual objects when the procedure is finished, even though they will not want to trade complete bundles. However, this type of envy will be difficult to identify under the compensation procedure, because the discounts that players receive apply only to their complete bundles; they cannot be broken down into partial discounts and associated with individual objects. The compensation procedure thus establishes envy-freeness for the highest aggregate utility that any player can identify. In order to obtain an efficient *bundling* of objects, players may still have to resort to computational support in order to improve their best initial guess.

7 Conclusions

Our objective in this paper was to formally develop a practical procedure for multilateral problems of fair-division. Our procedural approach eliminates envy in a “natural” way: it first identifies the players who are non-envious (we showed that there will always be at least one if the assignment of bundles is efficient), and then it compensates those players whose maximum envy is directed towards non-envious players. By first beginning with an efficient assignment, we formulated the procedure for establishing envy-freeness using intuitive, plausible, and manageably simple steps. We then showed that the procedure is also capable of guiding parties to an efficient assignment when the bundles are fixed. Moreover, the outcome of the compensation procedure is the same no matter which efficient assignment is eventually reached.

In our analysis, we placed no bound on the level of individual compensa-

tions. Consequently, we do not preclude the possibility that an individual may end up being paid by the others to take a bundle of goods. In the context of fair division we do not find this problematic at all. Indeed, if a group does not wish to exclude any of its members, then there is no reason why the group should not subsidize a member for receiving an undesired bundle. Moreover, the qualification constraint guarantees that subsidization is never a consequence of a player's insufficient valuation of the complete set of objects to be distributed.

As with all cooperative procedures, our compensation scheme is theoretically vulnerable to strategic manipulation. However, practically speaking, this requires detailed information on the other players' preferences. If players' bids are disclosed simultaneously, strategic bids under incomplete information can easily backfire: an untruthful player may have to pay more than what she thinks her bundle is worth or may envy some other players for their bundles. By contrast, since a player's compensation is based on her own subjective assessments, truthful behavior will always guarantee her envy-freeness, regardless of the others' behavior. In practice, this insurance creates a disincentive for attempting to distort one's preferences.

Implemented as a computer algorithm, our compensation procedure is polynomially bounded, thus making it comparable to alternative algorithms. But this is not our point: as a *procedure*, it is not meant to be run on a machine; instead it is designed to be used live in a mediation process.

Appendix

Envy-free division of burdens

Consider the following example of four players that have to divide a number of chores among each other. The total compensation for all chores is: $-C = 300$. The utilitarian assignment bundles the chores and thereby determines the minimum total requested payment $-M > 0$. Table 5 shows players (negative) bids for the four bundles of burdens.¹⁷

Table 5. Players' requested payments

	\bar{B}_1	\bar{B}_2	\bar{B}_3	\bar{B}_4
P1	-50	-80	-90	-80
P2	-40	-60	-85	-90
P3	-100	-60	-75	-65
P4	-50	-65	-90	-70
Initial reward	50	60	75	70

¹⁷ The entries in Table 5 are equal to the entries in Table 1 minus 100.

Players' demands under the utilitarian assignment (along the diagonal) are satisfied through initial payments, which sum to 255. The total compensation of 300 thus leaves an extra 45 to be granted to the players in the form of bonuses. The assessment matrix can be computed by subtracting the diagonal entry from each column.

The resulting initial assessment matrix is identical to the one in Table 2 where row i now shows Player i 's assessment of Player j 's bundle of chores. Consequently, the remaining surplus is also the same as in the example of Sect. 3. Hence, the compensation procedure together with the distribution of the remaining surplus yields the same outcome, whether for the distribution of goods with a total contribution, or for the distribution of chores with a total compensation, or for a combination of goods and chores with a total cost or compensation.

Numerical calculations for the average discount method

In what follows, (i)–(iv) denote the steps of the average discount method algorithm.

Beginning with Player 1, the procedure places this player in set D (i). Since Player 2 sees her bundle tied with that of Player 1 (see Row 2 of Table 4), Player 2 is added to set D (ii). In further comparisons, one now finds that Players 3 and 4 see their bundles tied with that of Player 2, so they must be added to D as well (ii). With all players included in D , the remaining surplus is divided equally among them (iii). The updated discounts are (5, 15, 15, 10), and the procedure for Player 1 ends because the surplus is used up (iv).

Beginning with Player 2 (i), Players 3 and 4 must be added to set D (ii). With only these players included in D , each one's discount can be raised by 5 before Player 1 experiences envy (see Row 1 of Table 4). This leaves a remaining surplus of $20 - 15 = 5$ (iii), so the procedure continues (iv). Player 1 now sees her bundle tied with those of Players 3 and 4, so she must be added to D (ii). The remaining surplus is divided equally among all players giving each an additional 1.25 (iii). The discounts are (1.25, 16.25, 16.25, 11.25), and the procedure for Player 2 ends (iv).

Beginning with Player 3 in set D (i), there is no player who feels tied with this player (ii). Player 3's discount can thus be raised by 5 without making any other player envious (iii). With a remaining surplus of $20 - 5 = 15$, the procedure continues (iv). Now Player 1 is tied with Player 3, Player 2 is tied with Player 1 and Player 4 is tied with Player 2, so they all must be added to D (ii). The remaining surplus is divided equally, giving each player an additional 3.75 (iii). The discounts are now (3.75, 13.75, 18.75, 8.75), and the procedure for Player 3 ends (iv).

Beginning with Player 4 in set D (i), Player 3 must be added because of the tie with Player 4 (ii). The discounts of these two players can be raised by 5 before Player 1 becomes envious (iii). With a remaining surplus of $20 - 10 = 10$, the procedure continues (iv). Now Player 1 is added to D and, therefore, also Player 2, because of the tie with Player 1 (ii). The remaining

Table 6. Players' bids before compensations are made

	\bar{B}_4	\bar{B}_1	\bar{B}_3	\bar{B}_2
P1	20	50	10	20
P2	10	60	15	40
P3	35	0	25	40
P4	30	50	10	35
	0	0	0	0





Table 7. Players' valuations of bundles after the first round of compensations

	\bar{B}_4	\bar{B}_1	\bar{B}_3	\bar{B}_2
P1	50	50	10	35
P2	40	60	15	55
P3	65	0	25	55
P4	60	50	10	50
	30	0	0	15



surplus is divided equally among all players, giving each an additional 2.5 (iii). The discounts are (2.5, 12.5, 17.5, 12.5), and the procedure for Player 4 ends (iv).

Eliminating envy-cycles with the compensation procedure

The following example shows how our modified compensation procedure with ex-post equal payments works for a non-utilitarian assignment. Consider the bid matrix shown in Table 6, which reproduces Table 1 with columns exchanged. The framed bids indicate the non-utilitarian assignment. The corresponding envy graph, constructed on the basis of bids, is denoted by \tilde{G} .

Strictly following the compensation procedure, Players 1 and 4 are now compensated by 30 and 15. (Player 3 is not yet compensated, because she envies an envious player.) Accordingly, the bid matrix changes to Table 7.

Now Players 3 and 4 envy Player 1. In order to eliminate their envy, they receive compensations of 40 and 10, respectively. The result is shown in Table 8. However, this second round of compensations involving the last envious players (3 and 4) creates new envy: Player 4 feels tied with 1 and Player 1 feels tied with 2, but Player 2 now envies 4, thus creating a cycle in the directed graph \tilde{G} .

The modified procedure now calls for a trade in the opposite direction of the arrows in \tilde{G} . Hence, Player 4 gives her bundle to Player 2, Player 2 gives her bundle to Player 1, and Player 1 gives her bundle to Player 2. In addi-

Table 8. Players' valuations of bundles after the second round of compensations

	\bar{B}_4	\bar{B}_1	\bar{B}_3	\bar{B}_2	
P1	50	50	50	45	Graph \tilde{G} $3 \Rightarrow 1 \Rightarrow 2 \rightarrow 4$ $4 \Rightarrow$
P2	40	60	55	65	
P3	65	0	65	65	
P4	60	50	50	60	
	30	0	40	25	

tion all compensations must be returned, and the compensation procedure is restarted. Note that the bid matrix after the required permutation of columns is now given by Table 1. As we know this new assignment is utilitarian and the compensation procedure with ex-post payments will lead to an envy-free outcome.

References

Alkan A, Demange G, Gale D (1991) Fair allocation of indivisible goods and criteria of justice. *Econometrica* 59: 1023–1039

Aragones E (1995) A derivation of the money Rawlsian solution. *Soc Choice Welfare* 12: 267–276

Brams SJ, Kilgour DM (2001) Competitive fair division. *J Political Econ* (forthcoming)

Brams SJ, Taylor AD (1996) Fair division – from cake-cutting to dispute resolution. Cambridge University Press, Cambridge, MA

Foley DK (1967) Resource allocation and the public sector. *Yale Econ Essays* 7: 45–98

Klijn F (2000) An algorithm for envy-free allocations in an economy with indivisible objects and money. *Soc Choice Welfare* 17: 201–216

Knaster B (1946) Sur le problème du partage pragmatique de H Steinhaus. *Ann Soc Polonaise Math* 19: 228–230

Steinhaus H (1948) The problem of fair division. *Econometrica* 16: 101–104

Su FE (1999) Rental harmony: Sperner’s lemma in fair division. *Am Math Monthly* 106: 930–942

Svensson LG (1983) Large indivisibles: An analysis with respect to price equilibrium and fairness. *Econometrica* 51: 939–954

Tadenuma K, Thomson W (1991) No-envy and consistency in economies with indivisible goods. *Econometrica* 59: 1755–1767