A Dual Approach to Fair Division of Indivisible Goods

Maciej Bartkowiak

June 2026

Abstract

We study a simple fair-division procedure for allocating indivisible goods among agents with identical, additive, and strictly ordered valuations. The target fairness notion is envy-freeness up to any good (EFX), a strong relaxation of envy-freeness for indivisible goods (Caragiannis et al., 2019; Plaut and Roughgarden, 2020). Instead of constructing the allocation by assigning goods one by one, the proposed procedure starts from a dual allocation in which every agent initially holds every good. The algorithm then repeatedly deletes a duplicated good from an agent with a currently most valuable bundle, always choosing the most valuable duplicated good available in that bundle. We prove that the procedure terminates with a feasible allocation and that the resulting allocation is EFX. The note is intended as a short preliminary conference version; its main contribution is the compact dual perspective and the accompanying proof for this restricted valuation setting.

1 Introduction

Fair division studies how to allocate resources among agents while respecting a chosen notion of fairness. Its classical origins go back to the formal study of division problems by Steinhaus (1948), and envy-freeness has long been one of the central fairness requirements (Varian, 1974). When the resources are indivisible goods, exact envy-freeness can be too demanding or impossible to satisfy, so the literature often considers approximate envy-freeness notions for discrete allocations (Amanatidis et al., 2022; Bouveret and Lang, 2008).

One influential relaxation is envy-freeness up to any good (EFX). Informally, an allocation is EFX if every agent weakly prefers her own bundle to any other agent's bundle after removing any single good from that other bundle. EFX is stronger than the more common "up to one good" relaxation and remains a central object in the fair allocation of indivisible goods (Caragiannis et al., 2019; Plaut and Roughgarden, 2020). General existence and computation questions for EFX are delicate, but several positive results are known in restricted settings, for example for three agents or for particular graph-structured instances (Chaudhury et al., 2020; Christodoulou et al., 2023).

This note considers a very restricted, but transparent, setting. Let $[m] = \{1, \ldots, m\}$ be the set of goods and let $[n] = \{1, \ldots, n\}$ be the set of agents. For a subset of goods $S \subseteq [m]$ and an integer t, let $\Pi_t(S)$ denote the set of all t-partitions of S. An allocation is an n-partition

$$(A_1, A_2, \ldots, A_n) \in \Pi_n([m]),$$

where A_i is the bundle allocated to agent i.

The preference of each agent over bundles is described by a valuation function. Throughout the note, valuations are assumed to be non-negative, identical, strictly ordered on single goods, and additive. Thus there is one valuation v such that $v = v_i = v_j$ for all $i, j \in [n]$, the goods are indexed so that

$$0 \leq v(1) < v(2) < \cdots < v(m),$$

and for every $S \subseteq [m]$,

$$v(S) = \sum_{g \in S} v(g).$$

An allocation $X = (X_1, \ldots, X_n)$ is called

- **envy-free (EF)** if $v(X_i) \geq v(X_j)$ for all $i, j \in [n]$;
- **envy-free up to any good (EFX)** if

  $$v(X_i) \geq \max_{g \in X_j} v(X_j \setminus \{g\})$$

  for all $i, j \in [n]$, with the comparison for an empty X_j understood as vacuous.

For later use, define the least valuable good of a nonempty bundle S

by

$$g^-(S) = \arg\min_{g \in S} v(g),$$

and define the value of S after deleting its least valuable good by

$$v^-(S) = v\big(S \setminus \{g^-(S)\}\big) = v(S) - v(g^-(S)).$$

For the empty bundle we set $v^-(\emptyset) = 0$. Since the valuation is additive and the goods are strictly ordered, deleting the least valuable good leaves the largest possible remaining value. Therefore, under identical valuations, EFX is equivalently the condition

$$v(X_i) \geq v^-(X_j) \qquad \text{for all } i, j \in [n].$$

The contribution of the note is a short dual algorithm and proof. Rather than assigning goods directly, the algorithm begins from the infeasible dual allocation in which each agent holds every good. It then makes the allocation feasible by deleting duplicated occurrences. The rule is greedy in the dual space: delete the most valuable duplicated good from an agent whose current bundle has maximum value. The theorem below shows that this simple deletion rule returns an EFX allocation in the identical-additive setting.

2 Dual approach and the algorithm

For the purpose of the Dual Fair Division Instance (DFDI), start from the forbidden dual allocation in which every good is held once by every agent:

$$DA = \{A_i : i \in [n]\}, \qquad A_i = [m] \text{ for every } i \in [n].$$

The procedure obtains a feasible allocation by deleting duplicated occurrences of goods. For monitoring the progress of the procedure, let

$$r_g = |\{i \in [n] : g \in A_i\}|$$

be the number of current occurrences of good g. Initially $r = (n, \ldots, n)$, and a feasible allocation is reached when $r = (1, \ldots, 1)$. At each step the algorithm chooses an agent with a currently most valuable bundle and strikes from that bundle the most valuable good that is still duplicated.

Algorithm 1: DFDI to FDI algorithm

Input: Agents $[n]$, goods $[m]$, identical additive valuation v, and the dual allocation $A_i \leftarrow [m]$ for every $i \in [n]$

Output: A feasible allocation $A = (A_1, \ldots, A_n)$ in which every good is assigned exactly once

```
Initialize r_g ← n for every g ∈ [m], and N ← [n];
while there exists g ∈ [m] with r_g > 1 do
    choose i ∈ N such that v(A_i) = max_{k∈N} v(A_k);
    S_i ← {g ∈ A_i : r_g > 1};
    if S_i ≠ ∅ then
        g* ← arg max_{g∈S_i} v(g);
        A_i ← A_i \ {g*};
        r_{g*} ← r_{g*} − 1;
    else
        N ← N \ {i};
return A;
```

Theorem 2.1 (EFX guarantee). *Assume identical additive valuations with* $0 \leq v(1) < \cdots < v(m)$*. Algorithm 1 terminates with a feasible allocation, and the returned allocation is EFX.*

Proof. The algorithm terminates because every deletion step decreases $\sum_{g \in [m]} r_g$ by one, so there can be at most $m(n-1)$ deletion steps. The other kind of step removes an agent from N, and this can happen at most n times. Moreover, an agent removed from N can never again have a duplicated good in her bundle, since the numbers r_g only decrease. Thus, if some duplicated good remained, at least one active agent would still hold a duplicated occurrence. Hence, when the algorithm stops, every good has exactly one remaining occurrence, so the output is a feasible allocation.

Let $V = v([m])$. During the algorithm, write

$$D_i = [m] \setminus A_i$$

for the set of goods already struck from agent i's dual bundle. Then

$$v(A_i) = V - v(D_i).$$

Thus choosing an agent with maximum current bundle value is the same as choosing, among active agents, an agent with minimum deleted value $v(D_i)$.

We prove EFX in the form $v(A_i) \geq v^-(A_j)$ for all ordered pairs of agents i, j. Fix $i, j \in [n]$. If $A_j = \emptyset$, then $v^-(A_j) = 0$, so the claim follows from non-negativity. If $i = j$, then $v(A_j) \geq v^-(A_j)$, again by non-negativity. Hence assume $i \neq j$ and $A_j \neq \emptyset$.

Let

$$g = g^-(A_j)$$

be the least valuable good in agent j's final bundle. Since the final allocation is feasible and $i \neq j$, good g is not in A_i at the end. Therefore there is a step at which the algorithm strikes g from agent i. Look at the moment immediately before this step, and denote the deleted sets at that moment by D_i^τ and D_j^τ.

At that moment, agent j is still active: indeed, j still holds g, agent i also still holds g just before it is struck, and therefore $r_g > 1$. Since the algorithm chooses i with maximum current bundle value among active agents, we have

$$v(A_i^\tau) \geq v(A_j^\tau).$$

Using $v(A_\ell^\tau) = V - v(D_\ell^\tau)$, this is equivalent to

$$v(D_i^\tau) \leq v(D_j^\tau). \tag{1}$$

Now consider what can happen after this moment. When g is struck from i, it is the most valuable currently duplicated good in A_i^τ. Therefore no good more valuable than g can be struck from i later: such a good either was already struck from i, or was already unique at time τ, and the repetition counts r_h never increase. Hence every good struck from i after g has value strictly smaller than $v(g)$.

On the other hand, because g is the least valuable good that remains in A_j at the end, every good h with $v(h) < v(g)$ is eventually struck from j. None of these smaller goods could have been struck from j before time τ, because before time τ the higher-valued good g was still in A_j and still duplicated, so the algorithm would strike a good of value at least $v(g)$ from j before striking any such h. Consequently, every possible later strike from i after g is matched by a strike of the same good from j after time τ.

Let

$$L_i = D_i \setminus \big(D_i^\tau \cup \{g\}\big) \qquad \text{and} \qquad L_j = D_j \setminus D_j^\tau$$

be the goods struck later from i after the strike of g, and the goods struck later from j after time τ, respectively. The previous paragraph gives $L_i \subseteq L_j$, and hence $v(L_i) \leq v(L_j)$. Combining this with (1),

$$\begin{aligned} v(D_i) &= v(D_i^\tau) + v(g) + v(L_i) \\ &\leq v(D_j^\tau) + v(g) + v(L_j) \\ &= v(D_j) + v(g). \end{aligned}$$

Finally, translating back from deleted bundles to final bundles gives

$$\begin{aligned} v(A_i) &= V - v(D_i) \\ &\geq V - v(D_j) - v(g) \\ &= v(A_j) - v(g) \\ &= v\big(A_j \setminus \{g^-(A_j)\}\big) \\ &= v^-(A_j). \end{aligned}$$

Since i and j were arbitrary, the returned allocation is EFX. $\square$

3 Conclusion

This preliminary note presents a dual deletion view of a simple EFX allocation procedure for identical, additive, and strictly ordered valuations. Starting from the infeasible allocation in which every agent holds every good, the algorithm removes duplicated goods from currently largest bundles until each good remains exactly once. The proof shows that the final allocation satisfies the equivalent condition $v(A_i) \geq v^-(A_j)$ for every pair of agents. Future work should clarify whether the same dual perspective can be extended beyond identical valuations, beyond strict single-good orderings, or to stronger efficiency requirements.

References

Amanatidis, G., Aziz, H., Birmpas, G., Filos-Ratsikas, A., Li, B., Moulin, H., et al. (2022). Fair division of indivisible goods: A survey.

Bouveret, S. and Lang, J. (2008). Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *Journal of Artificial Intelligence Research*, 32.

Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A. D., Shah, N., and Wang, J. (2019). The unreasonable fairness of maximum nash welfare. *ACM Transactions on Economics and Computation*, 7(3).

Chaudhury, B. R., Garg, J., and Mehlhorn, K. (2020). EFX exists for three agents. *CoRR*, abs/2002.05119.

Christodoulou, G., Fiat, A., Koutsoupias, E., and Sgouritsa, A. (2023). Fair allocation in graphs. In *Proceedings of the 24th ACM Conference on Economics and Computation*, EC '23, page 473–488, New York, NY, USA. Association for Computing Machinery.

Plaut, B. and Roughgarden, T. (2020). Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics*, 34(2):1039–1068.

Steinhaus, H. (1948). The problem of fair division. *Econometrica*, 16(1):101–104.

Varian, H. R. (1974). Equity, envy and efficiency. *Journal of Economic Theory*, 9:63–69.