Preprint – IJCAI 2025: This is the accepted version made available for conference attendees.
Do not cite. The final version will appear in the IJCAI 2025 proceedings.

Dividing Conflicting Items Fairly

Ayumi Igarashi[1] , **Pasin Manurangsi**[2] , **Hirotaka Yoneda**[1]
[1]University of Tokyo
[2]Google Research
igarashi@mist.i.u-tokyo.ac.jp, pasin@google.com, yoneda-h@g.ecc.u-tokyo.ac.jp

Abstract

We study the allocation of indivisible goods under conflicting constraints, represented by a graph. In this framework, vertices correspond to goods and edges correspond to conflicts between a pair of goods. Each agent is allocated an independent set in the graph. In a recent work of Kumar et al. (AAMAS, 2024), it was shown that a maximal EF1 allocation exists for interval graphs and two agents with monotone valuations. We significantly extend this result by establishing that a maximal EF1 allocation exists for *any graph* when the two agents have monotone valuations. To compute such an allocation, we present a polynomial-time algorithm for additive valuations, as well as a pseudo-polynomial time algorithm for monotone valuations. Moreover, we complement our findings by providing a counterexample demonstrating a maximal EF1 allocation may not exist for three agents with monotone valuations; further, we establish NP-hardness of determining the existence of such allocations for every fixed number $n \geq 3$ of agents. All of our results for goods also apply to the allocation of chores.

1 Introduction

How can we allocate a resource fairly? This problem was first formalized by the pioneering work of [Steinhaus, 1949] and has since been extensively studied in the fields of economics, mathematics, and computer science under the umbrella of *fair division*. Applications of fair division arise in many real-life situations, including the allocation of courses among students [Budish *et al.*, 2017], the division of family inheritance among family members [Goldman and Procaccia, 2014], and the division of household chores between couples [Igarashi and Yokoyama, 2023].

A central notion of fairness in fair division is *envy-freeness*, which requires that every agent is allocated their most preferred bundle in the allocation. However, such a fairness guarantee is impossible to achieve when dealing with indivisible resources, such as courses, houses, or tasks. Consequently, recent literature on discrete fair division has focused on approximate fairness, exploring various concepts and algorithmic results [Amanatidis *et al.*, 2023]. One particular influential relaxation of envy-freeness is, *envy-freeness up to one good (EF1)*, introduced by [Budish, 2011], allowing agents to remove one good from others' bundle to eliminate envy. This concept has garnered significant attention over the past decade. It is known that for general classes of monotone valuations, an EF1 allocation exists and can be computed efficiently [Lipton *et al.*, 2004].

Most studies on fair division assume that any allocation is feasible. While this assumption may hold in some cases, many practical scenarios involve constraints that restrict the structure of allocations. For instance, consider the allocation of multiple offices among several people over different periods of time. If the time intervals associated with two offices overlap, they cannot be assigned to the same person. Similar constraints arise in job scheduling, where overlapping shifts cannot be allocated to the same employee. Another example is the allocation of players to sports teams. If two players have overlapping areas of expertise, it is preferable not to assign them to the same team. A versatile framework for modeling such constraints, explored in a series of recent papers [Chiarelli *et al.*, 2023; Hummel and Hetland, 2022; Kumar *et al.*, 2024], represents conflicts among indivisible resources using a graph structure, where vertices correspond to resources and edges represent conflicts.

Conflicting constraints introduce new challenges, as standard fairness and efficiency concepts often become unattainable. Notably, canonical efficiency concepts such as completeness and Pareto-optimality are incompatible with EF1 under these constraints. In fact, with conflicting constraints, a complete allocation that assigns all goods may not always exist. Furthermore, even if such an allocation exists, there are simple instances where a complete EF1 allocation is unattainable [Hummel and Hetland, 2022].

[Kumar *et al.*, 2024] studied chore allocation under conflicting constraints, observing that even on a path, EF1 is incompatible with Pareto-optimality for two agents with identical additive valuations.[1] To address this limitation, they introduced the notion of *maximality*. A maximal allocation ensures that no unassigned item can be feasibly allocated to some agent. Kumar et al. showed that for interval graphs—a common structure in job scheduling—a maximal EF1 alloca-

[1]In [Kumar *et al.*, 2024], Pareto-optimality is defined with respect to maximal allocations.

tion among two agents always exists and can be efficiently found for monotone valuations. However, the existence of a maximal EF1 allocation for more general graph families remains unresolved, offering a rich avenue for further research.

Our contributions. We study the allocation of indivisible goods under conflicting constraints. Our goal is to identify conditions under which a maximal EF1 allocation exists and can be efficiently computed. Our main contributions are:

1. **Two-Agent-Case**: We significantly extend the result of [Kumar *et al.*, 2024] by establishing that a maximal EF1 allocation exists for *any graph* when the agents have monotone valuations. Further, we develop efficient algorithms for finding such allocations, including a polynomial-time algorithm for additive valuations and a pseudo-polynomial-time algorithm for monotone valuations. Note that the two-agent case is of particular importance in fair division, with various applications, including inheritance division, house-chore division, and divorce settlements [Brams and Fishburn, 2000; Brams *et al.*, 2014; Igarashi and Yokoyama, 2023].
2. **Three or More Agents**: We establish a sharp dichotomy from the two-agent case in terms of both existence and computational complexity. First, we provide an example where a maximal EF1 allocation fails to exist, even for three agents with monotone valuations. While [Hummel and Hetland, 2022] previously identified a counterexample for four agents, no such example was known for three agents. We also prove the NP-hardnesss of determining the existence of a maximal EF1 allocation for a fixed number $n \geq 3$ of agents with monotone valuations.
3. **Chore Allocation**: Finally, we consider the problem of chore allocation, where each agent has a monotone non-increasing valuation. We show that the existence of a maximal EF1 allocation under identical valuations directly translates from the goods case, establishing that all our results for goods hold for chores as well.

Related work. There is a growing body of research on fair division under constraints. For a comprehensive survey on the topic, see [Suksompong, 2021]. Conflicting constraints in the context of fair division were introduced by [Chiarelli *et al.*, 2023] and have been further explored by [Hummel and Hetland, 2022; Kumar *et al.*, 2024; Biswas *et al.*, 2023; Li *et al.*, 2021]. [Chiarelli *et al.*, 2023] explored different fairness objectives from ours, focusing on partial allocations that maximize the egalitarian social welfare—defined as the value of a bundle received by the worst-off agent. [Hummel and Hetland, 2022] studied complete allocations satisfying fairness criteria such as EF1 and MMS (maximin fair share). [Biswas *et al.*, 2023] generalized the model of [Hummel and Hetland, 2022], taking into account capacity of resources. [Li *et al.*, 2021] and [Kumar *et al.*, 2024] considered an interval scheduling problem, with the goal of achieving fairness concepts such as EF1 and MMS. While [Li *et al.*, 2021] focused on goods allocation with flexible intervals, [Kumar *et al.*, 2024] examined chore allocation.

A related type of constraints to conflicting constraints is the connectivity constraints of a graph [Bouveret *et al.*, 2017; Bilò *et al.*, 2022], where each agent receives a connected bundle of a graph. Note that while connectivity constraints imposed by a sparse graph such as a tree allow fewer feasible allocations, in our setting, sparsity implies greater flexibility, as it increases the number of feasible allocations.

2 Preliminaries

For any natural number $s \in \mathbb{N}$, let $[s] = \{1, 2, \ldots, s\}$.

Problem instance. We use $M = [m]$ to denote the set of *goods* and $N = [n]$ to denote the set of *agents*. Let $G = (M, E)$ denote an undirected graph, where each vertex corresponds to a good and each edge corresponds to a conflict. Each agent i has a *valuation function* $v_i : 2^M \to \mathbb{R}_+$; here, $\mathbb{R}_+$ is the set of non-negative reals. We assume that $v_i(\emptyset) = 0$. The valuation of a single good $g \in M$, $v_i(\{g\})$, is denoted by $v_i(g)$. An instance of our problem is given by the tuple $(N, M, \mathcal{V}, G)$ where $\mathcal{V} = (v_1, v_2, \ldots, v_n)$ denotes a *valuation profile*. We use $K_{s,t}$ to denote a complete bipartite graph with a bipartition in which one part contains s vertices and another part includes t vertices.

Valuation function. A valuation function v_i is *monotone non-decreasing* if $v_i(S) \leq v_i(T)$ for every $S \subseteq T \subseteq M$. Unless specified otherwise, we refer to such a function simply as *monotone*. It is *additive* if $v_i(S) = \sum_{g \in S} v_i(g)$ holds for every $S \subseteq M$ and $i \in N$. A valuation profile $\mathcal{V}$ is called *identical* if every agent $i \in N$ has the same valuation function; in this case, we denote this function by v. Let $T(m)$ denote the time to compute valuation $v_i(S)$ for a given $S \subseteq M$.

Allocation. An *allocation* is an ordered subpartition $\mathcal{A} = (A_1, \ldots, A_n)$ of M where for every pair of distinct agents $i, j \in N$, $A_i \cap A_j = \emptyset$, $\cup_{i \in N} A_i \subseteq M$, and for each $i \in N$, A_i is an *independent set* of G, namely, there is no pair of goods in A_i that are adjacent to each other. The subset A_i is called the *bundle* of agent i. An allocation is *complete* if all goods are allocated, i.e., $\cup_{i \in N} A_i = M$.

Fairness and efficiency notions. An allocation is *envy-free* (EF) if for every pair of agents $i, j \in N$, we have $v_i(A_i) \geq v_i(A_j)$ [Gamow and Stern, 1958; Foley, 1967]. It is *envy-free up to one good* (EF1) if for every pair of agents $i, j \in N$, either $A_j = \emptyset$, or there exists some good $g \in A_j$ such that $v_i(A_i) \geq v_i(A_j \setminus \{g\})$ [Budish, 2011; Lipton *et al.*, 2004].

As discussed in the Introduction, observe that a complete allocation may not always exist (e.g., consider a complete graph K_2 with one agent). Further, even when a complete allocation exists, a complete EF1 allocation may not exist: For instance, in a setting with n agents and a star where the center has a value of 0 and each of the $n + 1$ leaves has a value of 1, an agent receiving the center cannot receive any other good, while at least one agent receives two or more leaves.

Besides completeness, another commonly used notion of efficiency in fair division is *Pareto-optimality*. Similar to the chore setting, one can show that EF1 is incompatible with Pareto-optimality for goods instances. As in [Kumar *et al.*, 2024], we therefore consider a relaxed efficiency notion of *maximality*. An allocation $\mathcal{A}$ is *maximal* if for every agent $i \in N$ and every unallocated good $g \in M \setminus \bigcup_{i \in N} A_i$, g

is adjacent to some good in A_i. Our goal is to achieve an allocation that simultaneously satisfies maximality and EF1.

3 Two Agents

In this section, we consider the case of two agents.

3.1 Cut-and-Choose Protocol

When there are two agents, we can use a well-known strategy called **cut-and-choose protocol** [Brams and Taylor, 1996].

Theorem 1. *Suppose that a maximal EF1 allocation always exists and can be computed in τ time for instances with two agents and identical valuations. Then, a maximal EF1 allocation always exists and can be computed in $\tau + 2T(m)$ time for instances with two agents.*

Proof. Let (A_1, A_2) a maximal EF1 allocation in a hypothetical scenario that valuations of both agents are v_1. Then, agent 2 chooses a preferred bundle, leaving the reminder for 1. The resulting allocation is maximal and EF1. □

Now, the question is whether a maximal EF1 allocation exists for identical valuations. In subsections 3.2-3.5, we assume that the valuations are identical and monotone, and valuations for both agents are represented by $v(S)$ $(S \subseteq M)$.

3.2 Proof Strategy of Kumar et al.

To describe our proof strategy, let us (informally) review Kumar et al.'s proof of the existence of maximal EF1 allocation when G is a path. At a high level, the idea of the proof is to construct a chain of maximal allocations $\mathcal{A}^{(0)}, \dots, \mathcal{A}^{(m-1)}$, as illustrated in Figure 1, satisfying the following two properties: (i) adjacent allocations only "differ slightly" and (ii) $(A_1^{(0)}, A_2^{(0)}) = (A_2^{(m-1)}, A_1^{(m-1)})$. The latter implies that the signs of $v(A_1^{(0)}) - v(A_2^{(0)})$ and $v(A_1^{(m-1)}) - v(A_2^{(m-1)})$ are different. Therefore, there exists an i that the signs of $v(A_1^{(i)}) - v(A_2^{(i)})$ and $v(A_1^{(i+1)}) - v(A_2^{(i+1)})$ are different. At this point, they show that at least one of $\mathcal{A}^{(i)}$ or $\mathcal{A}^{(i+1)}$ must be EF1, which shows the existence of a maximal EF1 allocation. This concludes the proof overview of [Kumar *et al.*, 2024] for path graphs. They also extended this method to interval graphs, although the construction of the chain of maximal allocations becomes significantly more involved. Indeed, as explained and formalized below, our main contribution is to give a construction of a chain for any graph.

3.3 Useful Definitions and Lemmata

Our construction will require several generalizations of definitions and lemmata from [Kumar *et al.*, 2024]. We believe that these tools can be useful beyond the context of our work. Firstly, we use the following definition of "adjacency". Compared to [Kumar *et al.*, 2024], our definition (Definition 2) is more relaxed, in that it does not place any requirement on $|A'_1 \setminus A_1|$ and $|A_2 \setminus A'_2|$ whereas Kumar et al.'s enforces that these are at most one. Such a relaxation is crucial since our "chain" constructed below violates the aforementioned Kumar et al.'s condition.

Definition 2. *A pair $(\mathcal{A}, \mathcal{A}')$ of allocations is* ordered adjacent *if $|A_1 \setminus A'_1| \le 1$ and $|A'_2 \setminus A_2| \le 1$.*

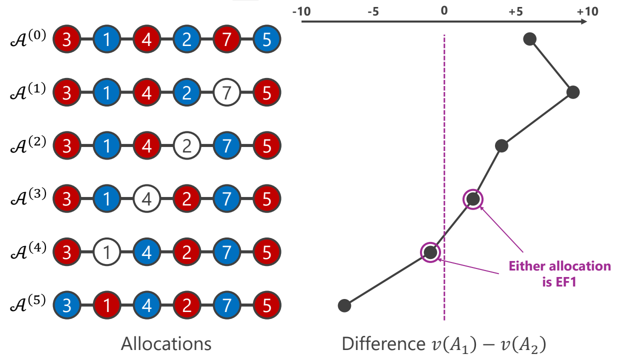

Figure 1: An example of chain of allocations for a path graph, $\mathcal{A}^{(0)}, \dots, \mathcal{A}^{(5)}$. The number written in each vertex is a valuation of the corresponding good. Red and blue vertices are those assigned to agent 1 and 2, respectively.

The following is the key lemma that enables to find an EF1 allocation at the point when $v(A_1) - v(A_2)$ crosses from positive to negative.

Lemma 3. *Let $(\mathcal{A}, \mathcal{A}')$ be any ordered adjacent pair of allocations. Further, assume that the following conditions hold:*

1. $v(A_1) \ge v(A_2)$
2. $v(A'_1) \le v(A'_2)$

Then, at least one of $\mathcal{A}$ or $\mathcal{A}'$ must be EF1.

Proof. Suppose that the allocation $\mathcal{A}$ is not EF1. Since $v(A_1) \ge v(A_2)$, agent a_2 envies agent a_1 even if one good is removed from A_1. Since $|A_1 \setminus A'_1| \le 1$, we must have $v(A_1 \cap A'_1) = v(A_1 \setminus (A_1 \setminus A'_1)) \ge v(A_2)$. As a result,

$$\begin{aligned} v(A'_1) &\ge v(A_1 \cap A'_1) \ge v(A_2) \\ &\ge v(A'_2 \cap A_2) = v(A'_2 \setminus (A'_2 \setminus A_2)), \end{aligned} \quad (1)$$

where the first and third inequalities hold because v is monotone. Now, consider the allocation $\mathcal{A}'$:

- Since $|A'_2 \setminus A_2| \le 1$, (1) implies that agent 1 does not envy agent 2 after removing a good from A'_2.
- Agent 2 does not envy agent 1 since $v(A'_1) \le v(A'_2)$.

Therefore, this allocation is EF1. □

We generalize Kumar et al.'s method by defining a **gapless chain**, and prove that any gapless chain must contain an EF1 allocation. We stress again that our requirements are weaker than the chain used in Kumar et al.'s[2], but still suffices to ensure EF1 for identical valuations.

Definition 4. *A sequence of allocation $\mathcal{A}^{(0)}, \dots, \mathcal{A}^{(k)}$ is a* **gapless chain** *if it satisfies the following conditions:*

1. $v(A_1^{(0)}) \ge v(A_2^{(0)})$.
2. $v(A_1^{(k)}) \le v(A_2^{(k)})$.

[2]Namely, we use a weaker adjacency notion and we only require the sign flip (first two conditions) whereas Kumar et al. require that $(A_1^{(0)}, A_2^{(0)}) = (A_2^{(k)}, A_1^{(k)})$.

3. *$(\mathcal{A}^{(i-1)}, \mathcal{A}^{(i)})$ is ordered adjacent for every $i \in [k]$.*

Lemma 5. *If $\mathcal{A}^{(0)}, \dots, \mathcal{A}^{(k)}$ is a gapless chain, there exists an $i \in \{0, \dots, k\}$ that $\mathcal{A}^{(i)}$ is EF1.*

Proof. From the first two conditions, there exists $i \in [k]$ that $v(A_1^{(i-1)}) \geq v(A_2^{(i-1)})$ and $v(A_1^{(i)}) \leq v(A_2^{(i)})$. Lemma 3 then implies that at least one of $\mathcal{A}^{(i-1)}$ or $\mathcal{A}^{(i)}$ is EF1. □

Given Lemma 5, our main task is thus to (efficiently) construct a gapless chain of maximal allocations (for any graph G). We devote the remainder of this section to this task.

3.4 Proof of Existence

Now, we prove that, a maximal EF1 allocation always exists for two agents, as stated below.

Theorem 6. *For $n = 2$ agents, any graph G and any monotone valuation v, there exists a maximal EF1 allocation.*

Although we only claim the existence in the above theorem, we will in fact present an algorithm for finding such an allocation, as its running time will be discussed in the next section. Our algorithm is presented in Algorithm 1 where the input S can be any maximal independent set of G. To prove Theorem 6, we use one that maximizes $v(S)$. In Figure 2, we provide an example of a chain constructed by the algorithm.

Algorithm 1 CHAINEF1$(S; G = (M, E), v)$

Require: $S = \{s_1, \dots, s_k\}$ is a maximal independent of G.
for $t \in M \setminus S$ **do**
  $\Gamma_t := \{i \in [k] \mid \{s_i, t\} \in E\}$
    ▷ Non-empty since S is a maximal independent set
  $p_t = \min_{i \in \Gamma_t} i$
  $q_t = \max_{i \in \Gamma_t} i$
$X_1 \leftarrow \emptyset$
for $t \in M \setminus S$ in increasing order of q_t **do**
  if t has no neighbor in X_1 **then**
    $X_1 \leftarrow X_1 \cup \{t\}$
$X_2 \leftarrow \emptyset$
for $t \in M \setminus S$ in decreasing order of p_t **do**
  if t has no neighbor in X_2 **then**
    $X_2 \leftarrow X_2 \cup \{t\}$
for $i = 0, \dots, k$ **do**
  $A_1^{(i)} = \{s_{i+1}, \dots, s_k\} \cup \{t \in X_1 \mid q_t \leq i\}$
  $A_2^{(i)} = \{s_1, \dots, s_i\} \cup \{t \in X_2 \mid p_t > i\}$
  if $\mathcal{A}^{(i)} = (A_1^{(i)}, A_2^{(i)})$ is EF1 **then**
    return $\mathcal{A}^{(i)}$
return NULL

We now prove a couple of crucial lemmata. Starting with the fact that each allocation is valid and maximal:

Lemma 7. *For every $i \in \{0, \dots, k\}$, $\mathcal{A}^{(i)}$ is a valid maximal allocation.*

Proof. First, we prove that the allocation is valid:

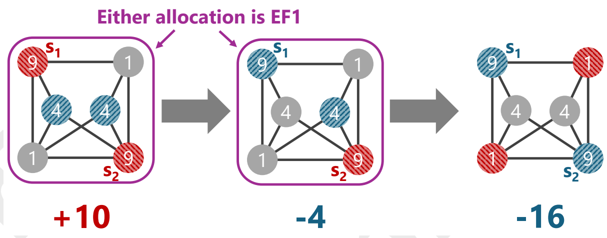

Figure 2: An example of a gapless chain of allocations, constructed by Algorithm 1. The number written in each vertex is a valuation of the corresponding good. Red and blue vertices are those assigned to agent 1 and 2, respectively.

- **No good is assigned to both agents.** This is obvious for goods in S. For good $t \in M \setminus S$, since $q_t \geq p_t$, at most one of the conditions $q_t \leq i$ or $p_t > i$ can hold; thus, it is assigned to at most one agent.
- **No adjacent goods are assigned to the same agent.** Consider two goods g, g' $(g \neq g')$ that are assigned to agent 1 and consider the following cases:
  1. Both goods are in S. They are not adjacent because S is an independent set.
  2. Both goods are in $M \setminus S$. They are not adjacent because $g, g' \in X_1$ and X_1 is an independent set.
  3. One good is from S and the other is from $M \setminus S$. Assume w.l.o.g. that $g \in S$ and $g' \in M \setminus S$. From this, we must have $g \in \{s_{i+1}, \dots, s_k\}$ and that $q_{g'} \leq i$. By the definition of $q_{g'}$, this ensures that g, g' are not adjacent.

Next, we prove maximality of $\mathcal{A}^{(i)}$, i.e., that no good in $M \setminus (A_1^{(i)} \cup A_2^{(i)})$ can be assigned to one of the agents. Let $g \in M \setminus (A_1^{(i)} \cup A_2^{(i)})$ be any unassigned good. Consider three following cases:

1. $q_g \leq i$. Since g is adjacent to $s_{q_g} \in A_2^{(i)}$, g cannot be assigned to agent 2. Moreover, since $g \notin A_1^{(i)}$, it must be that $g \notin X_1$. From how X_1 is constructed, there exists $t \in X_1$ such that $q_t \leq q_g$ and t is adjacent to g. As $t \in A_1^{(i)}$, g cannot be assigned to agent 1.
2. $p_g > i$. This case follows from an analogous argument to the first case.
3. $p_g \leq i < q_g$. In this case, g is adjacent to $s_{p_g} \in A_2^{(i)}$ and $s_{q_g} \in A_1^{(i)}$. Thus, t_j cannot be assigned. □

Next, we show that $\mathcal{A}^{(0)}, \dots, \mathcal{A}^{(k)}$ is a gapless chain. However, this requires an additional requirement that $v(S)$ is no smaller than $v(X_1), v(X_2)$.

Lemma 8. *If $v(S) \geq v(X_1)$ and $v(S) \geq v(X_2)$, then $\mathcal{A}^{(0)}, \dots, \mathcal{A}^{(k)}$ is a gapless chain.*

Proof. Note that $\mathcal{A}^{(0)} = (S, X_2)$ and $\mathcal{A}^{(k)} = (X_1, S)$. Thus, the first two conditions are satisfied by our assumption $v(S) \geq v(X_1), v(S) \geq v(X_2)$. Finally, for every $i \in [k]$, we can see that $A_1^{(i-1)} \setminus A_1^{(i)} = \{s_i\}$ and $A_2^{(i)} \setminus A_2^{(i-1)} = \{s_i\}$. Thus, $(\mathcal{A}^{(i-1)}, \mathcal{A}^{(i)})$ is ordered adjacent. □

Lemmas 5, 7 and 8 together immediately imply the following:

Lemma 9. *If $v(S) \geq v(X_1)$ and $v(S) \geq v(X_2)$, then Algorithm 1 outputs a maximal EF1 allocation.*

Our main theorem of this section (Theorem 6) now then follows easily from Lemma 9 by choosing an appropriate S.

Proof of Theorem 6. Let S be a maximal independent set of G that maximizes $v(S)$. Since X_1, X_2 are independent set and v is monotone, we have $v(S) \geq v(X_1)$ and $v(S) \geq v(X_2)$. Thus, Lemma 9 ensures that running Algorithm 1 on input S yields a maximal EF1 allocation. □

3.5 Algorithm

Recall in the proof of Theorem 6 that we pick S to be a maximal independent set with largest valuation. Doing so trivially would require enumerating through all 2^m subsets of M. In this section, we give a simple algorithm (Algorithm 2) that significantly improves upon this running time. In particular, it runs in polynomial-time for additive valuations and pseudo-polynomial time for general monotone valuations.

Algorithm 2 SWAPEF1($G = (M, E), v$)

```
g* ← argmax_{g∈M} v(g)
S^i ← any maximal independent set of G containing g*
i ← 0
while True do
    if CHAINEF1(S^i; G, v) ≠ NULL then
        return CHAINEF1(S^0; G, v)
    X_1^i, X_2^i ← X_1, X_2 in CHAINEF1(S^i; G, v)
    ℓ ← argmax_{ℓ'∈{1,2}} v(X_{ℓ'}^i)
    S^{i+1} ← any maximal independent set containing X_ℓ^i
    i ← i + 1
```

Our algorithm running time is stated below in Theorem 10. Note that in the second case, if the valuations $v(S)$ are all integers, the running time is pseudo-polynomial.

Theorem 10. *When there are $n = 2$ agents, there exists an algorithm that can find a maximal EF1 allocation and its running time is as follows:*

- $O(3^{m/3} \cdot (m \cdot T(m) + |E|))$ *for monotone valuations,*
- $O(B \cdot (m \cdot T(m) + |E|))$ *for monotone valuations such that the number of distinct values of $v(S)$ is at most B (i.e. $|\{v(S) \mid S \subseteq M\}| \leq B$), and,*
- $O(m \log m \cdot (m \cdot T(m) + |E|))$ *for additive valuations.*

Proof. Observe that the preprocessing time and the running time of each loop is $O(m \cdot T(m) + |E|)$. Thus, it suffices to bound the number of iterations of the loop in each case.

- For monotone valuations, by Lemma 9, each loop either terminates or we must have $v(X_\ell^i) > v(S^i)$. This implies $v(S^{i+1}) > v(S^i)$. This means that $S^0, S^1, \ldots$ are all distinct maximal independent sets of G. Since there are at most $3^{m/3}$ maximal independent sets [Moon and Moser, 1965], the number of iterations is at most $3^{m/3}$.
- Next, suppose that $v(S)$ can take at most B distinct values. From the argument above, $v(S^0), v(S^1), \ldots$ are distinct. Thus, the number of iterations is at most B.
- Finally, suppose that v is additive. We claim that the following holds for each loop i that does not terminate:

$$v(S^{i+1}) > \frac{m}{m-1} \cdot v(S^i). \qquad (2)$$

Before we prove (2), let us first use it to bound the number of iterations. Notice that (2) implies that, after loop i, we must have $v(S^{i+1}) > \left(\frac{m}{m-1}\right)^{i+1} \cdot v(S^0) \geq \left(\frac{m}{m-1}\right)^{i+1} \cdot v(g^*)$. Moreover, our choice of g^* ensures that $v(S^{i+1}) \leq m \cdot v(g^*)$. Thus, the number of iterations is at most $\log_{m/(m-1)} m + 1 = O(m \log m)$.

To see that (2) holds, first recall from Algorithm 1 that, if Algorithm 1 returns NULL, it has considered either (S^i, X_ℓ^i) or (X_ℓ^i, S^i) already and has determined that this is not EF1. Since $v(X_\ell^i) > v(S^i)$, this means that, for any good $g \in X_\ell^i$, we must have $v(S) < v(X_\ell^i \setminus \{g\})$. When we pick $g \in X_\ell^i$ with the largest $v(g)$, we have

$$v(g) \geq \frac{1}{|X_\ell^i|} v(X_\ell^i) \geq \frac{1}{m} v(X_\ell^i).$$

Hence, $v(S^i) < v(X_\ell^i \setminus \{g\}) \leq \frac{m-1}{m} v(X_\ell^i)$, proving (2). □

4 Three or More Agents

Negative Examples. While a maximal EF1 allocation always exists when there are two agents, it turns out that this is not the case for three agents or more.

Theorem 11. *For $n = 3$ agents, there exists an instance with identical monotone valuations where no maximal EF1 allocation exists.*

Proof Sketch. Consider the following instance with 7 goods. The graph consists of a complete bipartite graph $K_{3,3}$ with bipartition $X = \{1, 2, 3\}$ and $Y = \{4, 5, 6\}$ together with two edges $\{1, 7\}$ and $\{4, 7\}$. Each of the three agents has an identical monotone valuation v such that

- $v(\emptyset) = 0$;
- $v(S) = 1$ if $|S| = 1$ and $S \in \{\{1\}, \{4\}\}$;
- $v(S) = 2$ if $|S| = 1$ and $S \notin \{\{1\}, \{4\}\}$;
- $v(S) = 3$ if S is $\{2, 7\}$, $\{3, 7\}$, $\{5, 7\}$, or $\{6, 7\}$;
- $v(S) = 4$ for any other $S \subseteq M$.

Note that the symmetry on the graph and valuation holds: swapping 2 and 3, swapping 5, 6, and swapping $\{1, 2, 3\}$ and $\{4, 5, 6\}$ all lead to the same instance, and swapping the bundles of agents does not change whether the allocation is EF1. For this reason, there are only 6 maximal allocations to consider (refer to Figure 3). All of them are not EF1, because:

- Allocations #1, #4, #5, #6: one agent takes one good but but there is another agent who takes three goods.

- Allocation #2: $v(\{5,7\}) = 3$ but $v(\{1,2\}), v(\{1,3\}), v(\{2,3\}) = 4$.
- Allocation #3: $v(4) = 1$ but $v(5) = v(7) = 2$. □

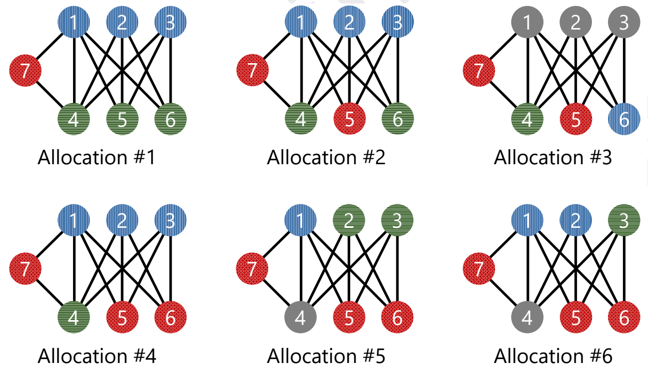

Figure 3: 6 maximal allocations to consider in the instance of Theorem 11. The vertices in red, blue, green, and gray are goods taken by agent 1, 2, 3, and no one, respectively.

The instance constructed in the proof of Theorem 11 uses an identical monotone valuation. It remains an open question for three agents with (even identical) additive valuations.

For $n \geq 4$ of agents, [Hummel and Hetland, 2022] presented an instance with identical additive valuations and a complete bipartite graph $K_{n-1,n-1}$ for which no complete EF1 allocation exists. In $K_{n-1,n-1}$, any maximal allocation is complete. Thus, this counterexample implies that achieving both EF1 and maximality is impossible for four agents with identical additive valuations. Here, we provide a smaller example using $K_{3,n-1}$, which turns out to be smallest since for $m \leq n+1$, there always exists a maximal EF1 allocation.

Proposition 12. *For every number $n \geq 4$ of agents, there is an instance with identical additive valuations, $m = n + 2$, and $G = K_{3,n-1}$ where no maximal EF1 allocation exists.*

Proposition 13. *For n agents with monotone valuations and $m \leq n + 1$ goods, there exists a maximal EF1 allocation.*

NP-Hardness. Next, we show that it is NP-hard to decide whether a maximal EF1 allocation exists, as stated below.

Theorem 14. *Given the graph and valuations, determining whether a maximal EF1 allocation exists is NP-hard for:*

1. *any fixed $n \geq 4$, even for identical and additive valuation, and,*
2. *$n = 3$, even for identical and monotone valuation.*

In fact, our proof can transform any negative example into an NP-hardness result, as stated more precisely below.

Lemma 15. *Suppose that there exists an instance $\tilde{I} = ([n], \tilde{M}, \tilde{v}, \tilde{G} = (\tilde{M}, \tilde{E}))$ (with identical valuation) where the number n of agents and the number $|\tilde{M}|$ of items are both constants, such that no maximal EF1 allocation exists. Then, it is NP-hard to decide whether a maximal EF1 allocation exists for n agents with identical valuations. Furthermore, if $\tilde{v}$ is additive, then this applies even for additive valuations.*

Theorem 14 is an immediate corollary of Lemma 15 where $\tilde{I}$ is the instance from Proposition 12 or Theorem 11. Note that we state Lemma 15 in this generic form so that, if subsequent work finds such an instance $\tilde{I}$ for additive valuation for $n = 3$, then the NP-hardness would follow as a corollary.

Reduction. We will reduce from the Independent Set (IS) problem, which is NP-hard [Karp, 1972]. In IS, we are given a graph $H = (V_H, E_H)$ and a positive integer t, and the goal is to decide whether H contains an IS of size t.

At a high-level, our reduction starts from $\tilde{I}$ and adds to it n copies of the graph H, where each good has the same value λ. Roughly speaking, we wish the i-th copy of H (denoted by X_i in the proof below) to give "extra goods" to the i-th agent, in case that agent envies some other agent by more than one good. The crux of the reduction is that such an agent can "catch up" (and thus satisfy EF1) iff there is a sufficiently large independent set in H. This is not yet a complete reduction since, H may not have a *maximal* independent set of a certain prescribed size. To alleviate this, we introduce "dummy goods" with zero value (denoted by Y_i below) to ensure that we can pick any desired number of goods from each copy of H. Finally, some additional edges are also added to ensure that each agent selects goods from a single copy of H.

For the proof below, we will use the following notation: for any valuation v and set S of goods, let $v^{-1}(S) := \min_{j \in S} v(S \setminus \{j\})$ denote the value of S after its most valuable good is removed. We use the convention $v^{-1}(\emptyset) = 0$.

Proof of Lemma 15. Recall $\tilde{I}$ from the lemma statement. Let[3] $\gamma := \min_{\tilde{A}} \max_{i,i' \in [n]} \left(\tilde{v}^{-1}(\tilde{A}_i) - \tilde{v}(\tilde{A}_{i'})\right)$ where the outer minimum is over all maximal allocation $\tilde{A}$ of $\tilde{I}$. By the assumption on $\tilde{I}$, we have $\gamma > 0$. Let $\lambda := \gamma/t$.

Let $(H = (V_H, E_H), t)$ denote the IS instance. Our reduction constructs the instance $I = ([n], M, v, G)$ as follows:

- **Goods**: $M = \tilde{M} \cup X_1 \cdots \cup X_n \cup Y_1 \cup \cdots \cup Y_n$ where $X_i = \{x_{i,w} \mid w \in V_H\}$ and $Y_i = \{y_{i,w} \mid w \in V_H\}$ are sets (each of size $|V_H|$) of additional goods.
- **Graph**: G contains the following edges:
  - (i) All edges in $\tilde{G}$,
  - (ii) $(x_{i,u}, x_{i,w})$ for all $i \in [n]$ and $(u, w) \in V_H$,
  - (iii) $(x_{i,w}, y_{i,w})$ for all $i \in [n]$ and $w \in V_H$,
  - (iv) all pairs of vertices in $(X_i \cup Y_i) \times (X_{i'} \cup Y_{i'})$ for all distinct $i, i' \in [n]$.
- **Valuation**: For all $S \subseteq M$, let $v(S) = \tilde{v}(S \cap \tilde{M}) + \lambda |S \cap X|$ where $X := X_1 \cup \cdots \cup X_n$. That is, the valuations on original goods remain the same, each good in X has value λ, and the goods in $Y_1 \cup \cdots \cup Y_n$ have value zero.

See Figure 4 for an illustration of the instance I.

It is clear that the reduction runs in polynomial time, and that, if $\tilde{v}$ is additive, then v is also additive.

[3] γ can be computed in $O(1)$ time by bruteforce.

(YES) Suppose that H contains an IS of size t. Let $\tilde{\mathcal{A}}^*$ be a maximal allocation of $\tilde{I}$ such that $\max_{i,i' \in [n]} \left(\tilde{v}^{-1}(\tilde{A}_i^*) - \tilde{v}(\tilde{A}_{i'}^*) \right) = \gamma$. Assume w.l.o.g. that $v^{-1}(\tilde{A}_1^*) \geq v^{-1}(\tilde{A}_2^*), \dots, v^{-1}(\tilde{A}_n^*)$. For each $i \in [n]$, we construct A_i^* as follows:

1. Let $c_i := \lceil \max\{0, \tilde{v}^{-1}(\tilde{A}_1^*) - \tilde{v}(\tilde{A}_i^*)\}/\lambda \rceil \leq t$.
2. Let S_i be any (non-necessarily maximal) IS of size c_i in H, which exists since H contains an IS of size t.
3. Let $A_i^* = \tilde{A}_i^* \cup \{x_{i,v}\}_{v \in S_i} \cup \{y_{i,v}\}_{v \in (V_H \setminus S_i)}$

Observe that each good in $X_i \cup Y_i$ can only belong to A_i^*, and there is no edge between goods in A_i^*. Thus, $\mathcal{A}^* = (A_1^*, \dots, A_n^*)$ is a valid allocation. To see that this is maximal, note that the goods from Y_i (resp., X_i) together with type-(iii) edges ensure that no other goods in X_i (resp., Y_i) can be added to A_i^*. Since at least one good from $X_i \cup Y_i$ is picked, type-(iv) edges ensure that no goods in $X_{i'} \cup Y_{i'}$ for $i' \neq i$ can be added to A_i^*.

Finally, we argue that $\mathcal{A}^*$ is EF1. To bound $v^{-1}(A_i^*)$, note that $v(A_i^*) = \tilde{v}(\tilde{A}_i^*) + c_i\lambda$. Consider two cases based on c_i.

- If $c_i = 0$, we have $v^{-1}(A_i^*) = \tilde{v}^{-1}(\tilde{A}_i^*) \leq \tilde{v}^{-1}(\tilde{A}_1^*)$.
- If $c_i > 0$, by definition of c_i, we have $v(A_i^*) < \tilde{v}^{-1}(\tilde{A}_1^*) + \lambda$. Thus, $v^{-1}(A_i^*) \leq v(A_i^*) - \lambda < \tilde{v}^{-1}(\tilde{A}_1^*)$.

Thus, in both cases, we have $v^{-1}(A_i^*) \leq \tilde{v}^{-1}(\tilde{A}_1^*)$.

On the other hand, for any $i' \in [n]$, the definition of $c_{i'}$ immediately implies $v(A_{i'}^*) \geq \tilde{v}^{-1}(\tilde{A}_1^*)$.

By the two previous paragraphs, $\mathcal{A}^*$ is EF1.

(NO) Suppose that H does not contain an IS of size t. Consider any maximal allocation $\mathcal{A} = (A_1, \dots, A_n)$ of I. Notice that the allocation $\tilde{\mathcal{A}} = (A_1 \cap \tilde{M}, \dots, A_n \cap \tilde{M})$ is maximal w.r.t. $\tilde{I}$. Thus, there exist $i, i' \in [n]$ such that $\tilde{v}^{-1}(\tilde{A}_i) - \tilde{v}(\tilde{A}_{i'}) \geq \gamma$. Due to type-(iv) edges, at most one of $A_{i'} \cap X_1, \dots, A_{i'} \cap X_n$ can be non-empty. Furthermore, type-(ii) edges imply that the non-empty set must correspond to an independent set in H. From our assumption, this implies $|A_{i'} \cap X| < t$. As a result,

$$\begin{aligned} v^{-1}(A_i) \geq \tilde{v}^{-1}(A_i) &\geq \gamma + \tilde{v}(\tilde{A}_{i'}) \\ &= \gamma + v(A_{i'}) - \lambda |A_{i'} \cap X| \\ &> \gamma + v(A_{i'}) - \lambda \cdot t \quad \overset{(\star)}{\geq} v(A_{i'}), \end{aligned}$$

where $(\star)$ is due to our choice of λ. Thus, $\tilde{A}$ is not EF1. □

5 Chore Allocation

In this section, we consider the chore version of our problem, where each agent i has a *monotone non-increasing* valuation function v_i. An allocation $\mathcal{A} = (A_1, \dots, A_n)$ is *envy-free up to one chore* (EF1 for chores) if for every pair of agents $i, j \in N$, either $A_i = \emptyset$, or $v_i(A_i \setminus \{c\}) \geq v_i(A_j)$ some $c \in A_i$ [Aziz *et al.*, 2022; Bhaskar *et al.*, 2021]. For identical valuations, the existence of a maximal EF1 allocation is equivalent for goods and chores: An allocation $\mathcal{A}$ is EF1 for chores under valuation v iff $\mathcal{A}$ is EF1 for goods under $-v$.

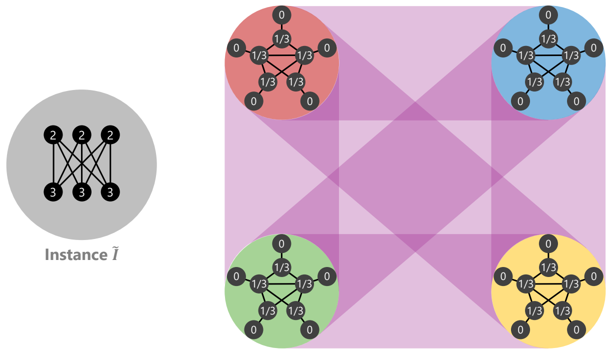

Figure 4: Instance I created with $n = 4$ instance given by Proposition 12 for $\tilde{I}$, and 5-vertex 7-edge graph for H, setting $t = 3$ (note that $\gamma = 1, \lambda = \frac{1}{3}$). The bands in purple represents type-(iv) edges.
.

Consider the case of two agents with monotone non-increasing valuations. Theorem 1 holds in this case, so we can assume w.l.o.g. that the agents have identical valuations v. By combining the discussion above with the fact that $-v$ is monotone non-decreasing, Theorem 6 guarantees the existence of a maximal EF1 allocation. Moreover, the algorithms presented in Theorem 10 remain applicable in this setting.

For three or more agents with monotone non-increasing valuations, there exist instances where a maximal allocation does not exist. In fact, the instances given in Theorem 12 and Proposition 11 use an identical valuation v, and the corresponding instance obtained by replacing v with $-v$ yields no maximal allocation that is EF1 for chores. Also, determining whether a maximal EF1 allocation exists is NP-hard also for monotone non-increasing valuations. This follows from Theorem 14 since the constructed instance uses an identical monotone non-decreasing valuation.

6 Conclusion

While we give a nearly complete picture of the existence for maximal EF1 allocations, there are still a few open questions left. First, for $n = 2$, our algorithm (Theorem 10) runs in pseudo-polynomial time for general monotone valuations. Is there a polynomial-time (in $m, T(m)$) algorithm for this task?

It might also be worthwhile considering special cases, by restricting either the valuations or the graphs. Examples are:

- Additive valuations: The existence of maximal EF1 allocation remains open only for the case $n = 3$ since our lower bound in Theorem 11 requires non-additive valuations, but the lower bound for $n \geq 4$ (Proposition 12) holds for additive (and identical) valuations.
- Uniform valuations: All $n \geq 3$ remains open in this case as we are not aware of any lower bound that holds for uniform valuations (namely, all goods are equally valued 1). We also note that the well-known Hajnal-Szemerédi theorem ([Hajnal and Szemerédi, 1970]) is equivalent to stating that a maximal EF1 allocation exists when the graph has maximum degree at most $n - 1$. Thus, a positive answer to the question for arbitrary graphs will significantly generalize the theorem.

Acknowledgements

This work was partially supported by JST FOREST Grant Numbers JPMJPR20C1. We thank the anonymous IJCAI 2025 for their valuable comments.

References

[Amanatidis *et al.*, 2023] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A Voudouris, and Xiaowei Wu. Fair division of indivisible goods: recent progress and open questions. *Artificial Intelligence*, page 103965, 2023.

[Aziz *et al.*, 2022] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. Fair allocation of indivisible goods and chores. *Autonomous Agents and Multi-Agent Systems*, 36(1):1–21, 2022.

[Bhaskar *et al.*, 2021] Umang Bhaskar, A. R. Sricharan, and Rohit Vaish. On Approximate Envy-Freeness for Indivisible Chores and Mixed Resources. In *Proceedings of Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2021)*, pages 1:1–1:23, 2021.

[Bilò *et al.*, 2022] Vittorio Bilò, Ioannis Caragiannis, Michele Flammini, Ayumi Igarashi, Gianpiero Monaco, Dominik Peters, Cosimo Vinci, and William S. Zwicker. Almost envy-free allocations with connected bundles. *Games and Economic Behavior*, 131:197–221, 2022.

[Biswas *et al.*, 2023] Arpita Biswas, Yiduo Ke, Samir Khuller, and Quanquan C. Liu. An algorithmic approach to address course enrollment challenges. In *Proceedings of 4th Symposium on Foundations of Responsible Computing (FORC 2023)*, pages 8:1–8:23, 2023.

[Bouveret *et al.*, 2017] Sylvain Bouveret, Katarína Cechlárová, Edith Elkind, Ayumi Igarashi, and Dominik Peters. Fair division of a graph. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 135–141, 2017.

[Brams and Fishburn, 2000] Steven J. Brams and Peter C. Fishburn. Fair division of indivisible items between two people with identical preferences: Envy-freeness, Pareto-optimality, and equity. *Social Choice and Welfare*, 17(2):247–267, 2000.

[Brams and Taylor, 1996] Steven J. Brams and Alan D. Taylor. *Fair division: from cake-cutting to dispute resolution*. Cambridge University Press, 1996.

[Brams *et al.*, 2014] Steven J. Brams, Marc Kilgour, and Christian Klamler. Two-person fair division of indivisible items: An efficient, envy-free algorithm. *Notices of the AMS*, 61(2):130–141, 2014.

[Budish *et al.*, 2017] Eric Budish, Gérard P. Cachon, Judd B. Kessler, and Abraham Othman. Course Match: a large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. *Operations Research*, 65(2):314–336, 2017.

[Budish, 2011] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

[Chiarelli *et al.*, 2023] Nina Chiarelli, Matjaz Krnc, Martin Milanic, Ulrich Pferschy, Nevena Pivac, and Joachim Schauer. Fair packing of independent sets. *Algorithmica*, 85(5):1459–1489, 2023.

[Foley, 1967] Duncan Foley. Resource allocation and the public sector. *Yale Economic Essays*, pages 45–98, 1967.

[Gamow and Stern, 1958] George Gamow and Marvin Stern. Puzzle-Math. *Viking Press*, 1958.

[Goldman and Procaccia, 2014] Jonathan Goldman and Ariel D. Procaccia. Spliddit: Unleashing fair division algorithms. *ACM SIGecom Exchanges*, 13(2):41–46, 2014.

[Hajnal and Szemerédi, 1970] András Hajnal and E. Szemerédi. Proof of a conjecture of P. Erdős. *Colloq Math Soc János Bolyai*, 4:601–623, 01 1970.

[Hummel and Hetland, 2022] Halvard Hummel and Magnus Lie Hetland. Fair allocation of conflicting items. *Autonomous Agents and Multi-Agent Systems*, 36(1):8, 2022.

[Igarashi and Yokoyama, 2023] Ayumi Igarashi and Tomohiko Yokoyama. Kajibuntan: a house chore division app. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, pages 16449–16451, 2023.

[Karp, 1972] Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller and James W. Thatcher, editors, *Proceedings of a symposium on the Complexity of Computer Computations, held March 20-22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, USA*, The IBM Research Symposia Series, pages 85–103, 1972.

[Kumar *et al.*, 2024] Yatharth Kumar, Sarfaraz Equbal, Rohit Gurjar, Swaprava Nath, and Rohit Vaish. Fair scheduling of indivisible chores. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 2345–2347, 2024.

[Li *et al.*, 2021] Bo Li, Minming Li, and Ruilong Zhang. Fair scheduling for time-dependent resources. In *Proceedings of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2021.

[Lipton *et al.*, 2004] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, pages 125–131, 2004.

[Moon and Moser, 1965] John W. Moon and Leo Moser. On cliques in graphs. *Israel Journal of Mathematics*, 3:23–28, 1965.

[Steinhaus, 1949] Hugo Steinhaus. Sur la division pragmatique. *Econometrica*, 17:315–319, 1949.

[Suksompong, 2021] Warut Suksompong. Constraints in fair division. *ACM SIGecom Exchanges*, 19(2):46–61, 2021.