Fair Division with Subsidy

Daniel Halpern and Nisarg Shah

University of Toronto, Canada

Abstract. When allocating a set of goods to a set of agents, a classic fairness notion called *envy-freeness* requires that no agent prefer the allocation of another agent to her own. When the goods are indivisible, this notion is impossible to guarantee, and prior work has focused on its relaxations. However, envy-freeness can be achieved if a third party is willing to subsidize by providing a small amount of money (divisible good), which can be allocated along with the indivisible goods.
In this paper, we study the amount of subsidy needed to achieve envy-freeness for agents with additive valuations, both for a given allocation of indivisible goods and when we can choose the allocation. In the former case, we provide a strongly polynomial time algorithm to minimize subsidy. In the latter case, we provide optimal constructive results for the special cases of binary and identical valuations, and make a conjecture in the general case. Our experiments using real data show that a small amount of subsidy is sufficient in practice.

Keywords: Fair division · Indivisible goods · Envy-freeness · Subsidy.

1 Introduction

How to fairly divide goods among people has been a subject of interest for millennia. However, formal foundations of fair division were laid less than a century ago with the work of Steinhaus [29], who proposed the cake-cutting setting where a *divisible* good is to be allocated to n agents with heterogeneous preferences. In the subsequent decades, allocation of divisible goods received significant attention [4, 16, 25, 32, 33]. When goods are divisible, one can provide strong fairness guarantees such as *envy-freeness* [17], which requires that no agent prefer the allocation of another agent to her own.

Most real-world applications of fair division, such as divorce settlement or inheritance division, often involve *indivisible* goods. In this case, envy-freeness is impossible to guarantee. For example, if the only available good is a ring, and two agents—Alice and Bob—want it, giving it to either agent would cause the other to envy. Recent research on fair allocation of indivisible goods has focused on achieving relaxed fairness guarantees [2, 11, 20, 27]. For example, *envy-freeness up to one good* requires that no agent prefer the allocation of another agent to her own after removing at most one good from the envied agent's bundle. This has lately been a subject of intensive research [7, 8, 26]. While giving the ring to Alice would satisfy this fairness guarantee, who can blame Bob for thinking that the allocation was unfair? After all, he received nothing!

Intuitively, it seems that if we have money at our disposal, it should help settle the differences and eliminate envy. But can it always help? Suppose that Alice values the ring at \$100 while Bob values it at \$150. If we give the ring to Alice, then Bob would require at least \$150 compensation to not envy Alice. But giving so much money to Bob would make Alice envy Bob. Upon some thought, it becomes clear that the only way to achieve envy-freeness is to give the ring to Bob and give Alice at least \$100 (but no more than \$150). Is this always possible? When can it be done?

In this paper, we study a setting where we allocate a set of indivisible goods along with some amount of a divisible good (a.k.a. money). The money can either be provided by a third party as a subsidy, or it could already be part of the set of goods available for allocation. Our primary research questions are:

> *Which allocations of indivisible goods allow elimination of envy using money? And how much money is required to achieve envy-freeness?*

1.1 Our Results

Suppose n agents have additive valuations (i.e., the value of a bundle is the sum of the values of the individual items) over m indivisible goods. Without loss of generality, we assume that the value of each agent for each good is in $[0, 1]$. We refer to an allocation of indivisible goods as *envy-freeable* if it is possible to eliminate envy by paying each agent some amount of money.

In Section 3, we characterize envy-freeable allocations and show how to efficiently compute the minimum payments to agents that are required to eliminate envy in a given envy-freeable allocation.

In Section 4, we study the size of the minimum subsidy (total payment to agents) required to achieve envy-freeness. When an (envy-freeable) allocation is given to us, we show that the minimum subsidy required is $\Theta(nm)$ in the worst case, even in the special cases of binary and identical valuations.

The picture gets more interesting when we are allowed to choose the allocation of indivisible goods. In this case, the minimum subsidy is at least $n - 1$ in the worst case. For the special cases of binary and identical valuations, we show that this optimal bound can be achieved through efficient algorithms. For general valuations, we show that it can be achieved for two agents, and conjecture this to be true for more than two agents.

Our experiments in Section 5 using synthetic and real data show that the minimum subsidy required in practice is much less than the worst-case bound.

1.2 Related Work

The use of money in fair allocation of indivisible goods has been well-explored. Much of the literature focuses on a setting where the number of goods is at most the number of agents. This is inspired from the classic rent division problem, where the goal is to allocate n indivisible goods to n agents and divide a total cost (rent) among the agents in an envy-free manner [30, 31]. In this case, Demange and Gale [13] show that the set of envy-free allocations have a lattice structure; we provide a similar result in Appendix A. Maskin [21] shows that envy-free allocations are guaranteed to exist given a sufficient amount of money; this is easy to show in our setting, so we focus on minimizing the amount of money required. Klijn [19] shows that envy-free allocations can be computed in polynomial time. Several papers focus on concepts other than (or stronger than) envy-freeness. For example, Quinzii [28] shows that the core coincides with competitive equilibria. Bikhchandani and Mamer [6] study the existence of competitive equilibria, which is a stronger requirement than envy-freeness. Ohseto [24] studies the existence of algorithms that are not only envy-free but also strategyproof. This restricted setting with *one good per agent* is substantially different from our general setting with potentially more goods than agents. Svensson [31] shows that in the restricted setting, envy-free allocations are automatically Pareto optimal. This is not true in our setting; and only a weaker condition is implied (Theorem 1).

Among the papers that consider more goods than agents, several consider settings which effectively reduce to one good per agent. For example, Haake et al. [18] consider a fixed partition of the goods into n bundles, so each bundle can be treated as a single good. In contrast, a large portion of our paper (Section 4.2) is devoted to finding the optimal bundling of goods. Further, they consider dividing a total cost of C among the agents, whereas we consider paying a non-negative amount of money to each agent. A natural reduction of our problem to their setting would set $C = 0$, compute the payments to agents (which could be negative), and increase all payments equally until they are non-negative. However, it is easy to check that under this reduction, our method requires less subsidy than theirs even for a fixed bundling, and significantly less if we optimize the bundling. Alkan et al. [1] allow more goods than agents, but add fictitious agents until the number of goods and agents are equal. As noted by Meertens et al. [22], their algorithm allocates at most one good to each real agent, throwing away the remaining goods (i.e. assigning them to fictitious agents).

Meertens et al. [22] study a setting more general than ours. They allow agents to have general preference relations over their allocated bundle of indivisible goods and amount of money. In this case, they show that envy-freeness and Pareto optimality may be incompatible regardless of the amount of money available. In contrast, in our setting with quasi-linear preferences, allocations that are both envy-free and Pareto optimal exist given a sufficient amount of money (see the discussion following Proposition 1). Beviá et al. [5] study a setting where each agent arrives at the market with a bundle of goods and an amount of money, and is interested in exchanging the goods and money with other agents. They assume that each agent brings at least as much money as her total value for the goods brought by all the agents, and induce budget-balanced transfers among the agents, making their results incomparable to ours.

To the best of our knowledge, no prior work studies the asymptotic amount of subsidy required to achieve envy-freeness, which is the focus of our work.

2 Preliminaries

For $k \in \mathbb{N}$, let $[k] = \{1, \ldots, k\}$. Let $\mathcal{N} = [n]$ denote the set of *agents*, and let $\mathcal{M}$ denote the set of m indivisible *goods*. Each agent i is endowed with a *valuation* function $v_i : 2^{\mathcal{M}} \to \mathbb{R}_{\geq 0}$ such that $v_i(\emptyset) = 0$. We assume that the valuation is *additive*: $\forall S \subseteq \mathcal{M}, v_i(S) = \sum_{g \in S} v_i(\{g\})$. To simplify notation, we write $v_i(g)$ instead of $v_i(\{g\})$. We denote the vector of valuations by $\mathbf{v} = (v_1, \ldots, v_n)$. We define an *allocation problem* to be the tuple $\mathcal{A} = (\mathcal{N}, \mathcal{M}, \mathbf{v})$.

For a set of goods $S \subseteq \mathcal{M}$ and $k \in \mathbb{N}$, let $\Pi_k(S)$ denote the set of ordered partitions of S into k bundles. Given an allocation problem $\mathcal{A}$, an *allocation* $\mathbf{A} = (A_1, \ldots, A_n) \in \Pi_n(\mathcal{M})$ is a partition of the goods into n bundles, where A_i is the bundle allocated to agent i. Under this allocation, the *utility* to agent i is $v_i(A_i)$, and the *utilitarian welfare* is $\sum_{i=1}^{n} v_i(A_i)$. The following fairness notion is central to our work.

Definition 1 (Envy-Freeness). *An allocation* $\mathbf{A}$ *is called envy-free (EF) if* $v_i(A_i) \geq v_i(A_j)$ *for all agents* $i, j \in \mathcal{N}$.

Envy-freeness requires that no agent prefer another agent's allocation over her own allocation. This cannot be guaranteed when goods are indivisible. Prior literature focuses on its relaxations, such as envy-freeness up to one good [10, 20], which can be guaranteed.

Definition 2 (Envy-Freeness up to One Good). *An allocation* $\mathbf{A}$ *is called* envy-free up to one good *(EF1) if, for all agents* $i, j \in \mathcal{N}$*, either* $v_i(A_i) \geq v_i(A_j)$ *or there exists* $g \in A_j$ *such that* $v_i(A_i) \geq v_i(A_j \setminus \{g\})$*. That is, it should be possible to remove envy between any two agents by removing a single good from the envied agent's bundle.*

We want to study whether (exact) envy-freeness can be achieved by additionally giving each agent some amount of a divisible good, which we refer to as *money*. We denote by $p_i \in \mathbb{R}$ the amount of money received by agent i, and by $\mathbf{p} = (p_1, \ldots, p_n)$ the vector of payments. Throughout most of the paper, we require that $p_i \geq 0$ for each agent i. This corresponds to the *subsidy model*, where a third party subsidizes the allocation problem by donating money. In Section 6, we discuss the implications of our results for other models of introducing monetary payments. One other obvious model is one in which there is no outside subsidy and envy is dealt with by agents paying each other. We show these models are essentially equivalent in the sense that any payments in one model can be translated to equivalent payments in the other. In our ring example, Bob giving Alice \$50 is equivalent to Alice receiving a \$100 subsidy with respect to relative utilities, which is all that matters for envy-freeness.

Given an allocation $\mathbf{A}$ and a payment vector $\mathbf{p}$, we refer to the tuple $(\mathbf{A}, \mathbf{p})$ as the *allocation with payments*. Under $(\mathbf{A}, \mathbf{p})$, the utility of agent i is $v_i(A_i) + p_i$. That is, agents have quasi-linear utilities (equivalently, they express their values for other goods with money as reference). With money, there is a common good to which agents can scale their utilities. Thus, unlike in settings without money, interpersonal comparisons of utilities make sense in our framework. Note that allocation $\mathbf{A}$ is equivalent to allocation with payments $(\mathbf{A}, \mathbf{0})$, where each agent receives zero payment. We can now extend the definition of envy-freeness to allocations with payments.

Definition 3 (Envy-Freeness). *An allocation with payments* $(\mathbf{A}, \mathbf{p})$ *is* envy-free *(EF) if* $v_i(A_i) + p_i \geq v_i(A_j) + p_j$ *for all agents* $i, j \in \mathcal{N}$*.*

We say that payment vector $\mathbf{p}$ is *envy-eliminating* for allocation $\mathbf{A}$ if $(\mathbf{A}, \mathbf{p})$ is envy-free. Let $\mathcal{P}(\mathbf{A})$ be the set of envy-eliminating payment vectors for $\mathbf{A}$.

Definition 4 (Envy-Freeable). *An allocation* $\mathbf{A}$ *is called* envy-freeable *if there exists a payment vector* $\mathbf{p}$ *such that* $(\mathbf{A}, \mathbf{p})$ *is envy-free, that is, if* $\mathcal{P}(\mathbf{A}) \neq \emptyset$*.*

Given an allocation problem $\mathcal{A}$, let $\mathcal{E}(\mathcal{A})$ denote the set of envy-freeable allocations. We drop $\mathcal{A}$ from the notation when it is clear from context.

Given an allocation $\mathbf{A}$, its *envy graph* $G_{\mathbf{A}}$ is the complete weighted directed graph in which each agent is a node, and for each $i, j \in \mathcal{N}$, edge (i, j) has weight $w(i, j) = v_i(A_j) - v_i(A_i)$. This is the amount of envy that agent i has for agent j, which can be negative if agent i strictly prefers her own allocation to the allocation of agent j. Note that by definition, $w(i, i) = 0$ for each $i \in \mathcal{N}$. A path P is a sequence of nodes $(i_1, \ldots, i_k)$, and its weight is $w(P) = \sum_{t=1}^{k-1} w(i_t, i_{t+1})$. The path is a *cycle* if $i_1 = i_k$. Given $i, j \in \mathcal{N}$, let $\ell(i, j)$ be the maximum weight of any path which starts at i and ends at j, and let $\ell(i) = \max_{j \in \mathcal{N}} \ell(i, j)$ be the maximum weight of any path starting at i.

3 Envy-Freeable Allocations

In this section, our goal is to characterize envy-freeable allocations of indivisible goods and, given an envy-freeable allocation, to find an envy-eliminating payment vector.

Looking more closely at $G_{\mathbf{A}}$, we can see that $\mathbf{A}$ being envy-free is equivalent to all edge weights of $G_{\mathbf{A}}$ being non-positive. We can extend this connection to the (potentially) larger set of envy-freeable allocations. Note that a permutation of $[n]$ is a bijection $\sigma : [n] \to [n]$.

Theorem 1. *For an allocation* $\mathbf{A}$*, the following statements are equivalent.*

(a) $\mathbf{A}$ *is envy-freeable.*
(b) $\mathbf{A}$ *maximizes the utilitarian welfare across all reassignments of its bundles to agents, that is, for every permutation* σ *of* $[n]$, $\sum_{i \in \mathcal{N}} v_i(A_i) \geq \sum_{i \in \mathcal{N}} v_i(A_{\sigma(i)})$*.*
(c) $G_{\mathbf{A}}$ *has no positive-weight cycles.*

Proof. We show $(a) \Rightarrow (b)$, $(b) \Rightarrow (c)$, and $(c) \Rightarrow (a)$.

$(a) \Rightarrow (b)$: Suppose $\mathbf{A}$ is envy-freeable. Then, there exists a payment vector $\mathbf{p}$ such that for all agents $i, j \in \mathcal{N}$, $v_i(A_i) + p_i \geq v_i(A_j) + p_j$, that is, $v_i(A_j) -$

$v_i(A_i) \le p_i - p_j$. Consider any permutation σ of $[n]$. Then, $\sum_{i \in \mathcal{N}} v_i(A_{\sigma(i)}) - v_i(A_i) \le \sum_{i \in \mathcal{N}} p_i - p_{\sigma(i)} = 0$.

(b) ⇒ (c): Suppose condition (b) holds. Consider a cycle $C = (i_1, \ldots, i_k)$ in $G_{\mathbf{A}}$. Consider the corresponding permutation σ_C under which $\sigma(i_t) = i_{t+1}$ for each $t \in [k-1]$, and $\sigma(i) = i$ for all $i \notin C$. Then,

$$
\begin{aligned}
w(C) &= \sum_{t=1}^{k-1} w(i_t, i_{t+1}) = \sum_{t=1}^{k-1} v_{i_t}(A_{i_{t+1}}) - v_{i_t}(A_{i_t}) \\
&= \sum_{t=1}^{k-1} \left(v_{i_t}(A_{i_{t+1}}) - v_{i_t}(A_{i_t})\right) + \sum_{i \notin C} (v_i(A_i) - v_i(A_i)) \\
&= \sum_{i \in \mathcal{N}} v_i(A_{\sigma(i)}) - v_i(A_i) \le 0.
\end{aligned}
$$

(c) ⇒ (a): Suppose $G_{\mathbf{A}}$ has no positive-weight cycles. Then, $\ell(i)$, which is the maximum weight of any path starting at i in $G_{\mathbf{A}}$, is well-defined and finite. Let $p_i = \ell(i)$ for each $i \in \mathcal{N}$. Note that $p_i \ge \ell(i,i) \ge w(i,i) = 0$ for each $i \in \mathcal{N}$. Hence, $\mathbf{p}$ is a valid payment vector. Also, by definition of longest paths, we have that for all $i, j \in \mathcal{N}$, $p_i = \ell(i) \ge \ell(j) + w(i,j) = p_j + v_i(A_j) - v_i(A_i)$. Hence, $(\mathbf{A}, \mathbf{p})$ is envy-free, and thus, $\mathbf{A}$ is envy-freeable. □

Theorem 1 provides a way to efficiently check if a given allocation $\mathbf{A}$ is envy-freeable. This can be done using the maximum weight bipartite matching algorithm [15] to check condition (b) or the Floyd-Warshall algorithm to check condition (c). The proof is provided in Appendix C.

Proposition 1. *Given an allocation* $\mathbf{A}$*, it is possible to check whether* $\mathbf{A}$ *is envy-freeable in* $O(mn + n^3)$ *time.*

Given Proposition 1, finding an envy-freeable allocation is easy: we can start from an arbitrary allocation $\mathbf{A}$ and use the maximum weight bipartite matching algorithm to find the reassignment of its bundles that maximizes utilitarian welfare, or we could simply compute the allocation that globally maximizes utilitarian welfare in $O(nm)$ time by assigning each good to the agent who values it the most.

But simply knowing an envy-freeable allocation $\mathbf{A}$ is not enough. We need to find a payment vector $\mathbf{p}$ such that $(\mathbf{A}, \mathbf{p})$ is envy-free. We would further like to minimize the subsidy required ($\sum_{i \in \mathcal{N}} p_i$). Such a payment vector can easily be computed in polynomial time through a linear program (see Appendix B.1). However, the next result shows that we can compute it in strongly polynomial time (polynomial in the number of inputs, rather than their size). In fact, this payment vector is precisely the one we constructed in the proof of Theorem 1.

Theorem 2. *For an envy-freeable allocation* $\mathbf{A}$*, let* $\mathbf{p}^*(\mathbf{A})$ *be given by* $p_i^*(\mathbf{A}) = \ell(i)$ *for all* $i \in \mathcal{N}$*, where* $\ell(i)$ *is the maximum weight of any path starting at* i *in* $G_{\mathbf{A}}$*. Then,* $\mathbf{p}^*(\mathbf{A}) \in \mathcal{P}(\mathbf{A})$*, and for every* $\mathbf{p} \in \mathcal{P}(\mathbf{A})$ *and* $i \in \mathcal{N}$*,* $p_i^*(\mathbf{A}) \le p_i$*. Further,* $\mathbf{p}^*(\mathbf{A})$ *can be computed in* $O(nm + n^3)$ *time.*

Proof. For simplicity, we denote $\mathbf{p}^*(\mathbf{A})$ as $\mathbf{p}^*$. When proving that condition (c) implies condition (a) in Theorem 1, we already showed that $\mathbf{p}^* \in \mathcal{P}(\mathbf{A})$. Thus, we simply need to argue that for every $\mathbf{p} \in \mathcal{P}(\mathbf{A})$, we have that $p_i^* \leq p_i$ for all $i \in \mathcal{N}$.

Fix $\mathbf{p} \in \mathcal{P}(\mathbf{A})$ and $i \in \mathcal{N}$. Consider the longest path starting at i in $G_{\mathbf{A}}$. Suppose it is $(i_1, \ldots, i_k)$. Hence, $i_1 = i$ and $w(i_1, \ldots, i_k) = \sum_{t=1}^{k-1} w(i_t, i_{t+1}) = p_i^*$. Because $(\mathbf{A}, \mathbf{p})$ is envy-free, we have that for each $t \in [k-1]$,

$$\begin{aligned} & v_{i_t}(A_{i_t}) + p_{i_t} \geq v_{i_t}(A_{i_{t+1}}) + p_{i_{t+1}} \\ & \Rightarrow p_{i_t} - p_{i_{t+1}} \geq v_{i_t}(A_{i_{t+1}}) - v_{i_t}(A_{i_t}) = w(i_t, i_{t+1}). \end{aligned}$$

Summing this over all $t \in [k-1]$, we get

$$p_{i_1} - p_{i_k} \geq w(i_1, \ldots, i_k) = p_i^* \Rightarrow p_i \geq p_i^* + p_{i_k} \geq p_i^*,$$

where the final transition holds because $i_1 = i$ and payments are non-negative.

Finally, $\mathbf{p}^*$ can be computed as follows. We first run the Floyd-Marshall (all-pairs shortest path) algorithm on the graph obtained by negating all edge weights in $G_{\mathbf{A}}$ to compute $\ell(i,j)$ for all $i, j \in \mathcal{N}$ in $O(nm + n^3)$ time. Then, we compute $\mathbf{p}^*$ in $O(n^2)$ time. □

We refer to $\mathbf{p}^*(\mathbf{A})$ as the *optimal payment vector* for $\mathbf{A}$. When clear from the context, we drop $\mathbf{A}$ from the notation.

We can also show that for an envy-freeable allocation $\mathbf{A}$, $\mathcal{P}(\mathbf{A})$ has a lattice structure and $\mathbf{p}^*$ is its unique minimum element; the proof is provided in Appendix A. In this lattice, the greatest lower bound (resp., the least upper bound) of two payment vectors is given by the coordinate-wise minimum (resp., maximum).

4 Minimizing and Bounding Subsidy

In this section, we investigate the minimum subsidy required to achieve envy-freeness. We are interested in both the computational complexity of computing the minimum subsidy required in a given allocation problem, and in the minimum subsidy required in the worst case over allocation problems. We consider cases where the (envy-freeable) allocation is given to us, and where we can choose such an allocation to minimize subsidy.

For an envy-freeable allocation $\mathbf{A}$, let $\text{sub}(\mathbf{A}) = \sum_{i \in \mathcal{N}} p_i^*(\mathbf{A})$ be the minimum subsidy required to make $\mathbf{A}$ envy-free. Then, in the former case, we want to compute $\sup_{\mathcal{A}} \max_{\mathbf{A} \in \mathcal{E}(\mathcal{A})} \text{sub}(\mathbf{A})$ and, in the latter case, we want to compute $\sup_{\mathcal{A}} \min_{\mathbf{A} \in \mathcal{E}(\mathcal{A})} \text{sub}(\mathbf{A})$.[1]

Without loss of generality, we assume that $v_i(g) \in [0,1]$ for each agent i and good g. If the valuations lie in $[0,T]$, the worst-case minimum subsidy and the bounds we provide would simply be multiplied by T, the largest value for any single good. We say that valuations are *binary* if $v_i(g) \in \{0,1\}$ for all agents i and goods g, and *identical* if $v_i(g) = v_j(g)$ for all agents i, j and goods g.

[1] Note that $\mathcal{E}(\mathcal{A}) \neq \emptyset$ because the allocation maximizing utilitarian welfare is always envy-freeable due to Theorem 1.

4.1 When the Allocation is Given

In cases where an envy-freeable allocation is already implemented, or if we desire to implement a specific allocation for reasons other than achieving envy-freeness, we may be given an allocation and asked to eliminate envy.

Theorem 2 already shows that we can efficiently compute the minimum amount of subsidy required. To study how much subsidy is needed in the worst case, we begin with the following simple observation.

Lemma 1. *For an envy-freeable allocation* $\mathbf{A}$*, no path in* $G_{\mathbf{A}}$ *has weight more than* m*.*

Proof. Since $G_{\mathbf{A}}$ has no positive-weight cycles, we only need to consider simple paths on which no agent appears twice. Consider a simple path $(i_1, \ldots, i_k)$. For $t \in [k-1]$, note that $w(i_t, i_{t+1}) = v_{i_t}(A_{i_{t+1}}) - v_{i_t}(A_{i_t}) \leq |A_{i_{t+1}}|$. Thus, the weight of the path is $\sum_{t=1}^{k-1} w(i_t, i_{t+1}) \leq \sum_{t=1}^{k-1} |A_{i_{t+1}}| = \left|\cup_{t=2}^{k} A_{i_t}\right| \leq m$, as desired. □

We can now pinpoint the subsidy required in the worst case. The upper bound uses Lemma 1 along with the fact that some agent must receive zero payment under the optimal payment vector.

Theorem 3. *When an envy-freeable allocation is given, the minimum subsidy required is* $(n-1)m$ *in the worst case.*

Proof. For the lower bound, consider the instance where $v_i(g) = 1$ for all agents i and goods g. Consider the allocation $\mathbf{A}$ which assigns all goods to a single agent i^*. It is easy to see that this is envy-freeable, and its optimal payment vector $\mathbf{p}$ has $p_i = m$ for $i \neq i^*$ and $p_{i^*} = 0$. Hence, we need $(n-1)m$ subsidy.

To prove the upper bound, note that the minimum subsidy required is the sum of weights of longest paths starting at different agents (Theorem 2). Using Lemma 1 and the fact that one agent must receive zero payment (otherwise all payments can be reduced while preserving envy-freeness, which would contradict the minimality of payments), this is at most $(n-1)m$. □

The lower bound uses an instance with identical binary valuations. Hence, Theorem 3 also holds for the special cases of binary and identical valuations.

4.2 When the Allocation Can Be Chosen

When we are allowed to choose the allocation, computing the minimum subsidy required is NP-hard. This is because checking whether zero subsidy is required is equivalent to checking whether an envy-free allocation exists, which is NP-hard even for identical valuations [9]. That said, it is possible to compute the minimum subsidy required using a simple integer linear program (see Appendix B.2).

Recall that when an envy-freeable allocation is *given*, in the worst case we need a subsidy of $(n-1)m$ (Theorem 3). But what if we were able to *choose* the allocation? We show that this does not help improve the bound by a factor larger than m.

Theorem 4. *When the allocation can be chosen, the minimum subsidy required is at least* $n-1$ *in the worst case, even in the special cases of binary valuations and identical valuations.*

Proof. Consider the instance with identical binary valuations where each agent values a special good at 1 and other goods at 0. Every allocation gives the special good to one of the agents. To achieve envy-freeness, each other agent must be paid at least 1. Hence, a subsidy of at least $n-1$ is needed. □

This raises a natural question: *Can we always find an envy-freeable allocation that requires a subsidy of at most* $n-1$*?* We answer this question affirmatively for the special cases of binary and identical valuations as well as any valuations with two agents. In addition, we make an interesting conjecture in the general case. First, we take a slight detour.

One promising approach to reducing the subsidy requirement is to start with an allocation that already has limited envy, for example, an allocation that is envy-free up to one good [10, 20]. For an envy-freeable EF1 allocation $\mathbf{A}$, each edge in $G_{\mathbf{A}}$ has weight at most 1, so each (simple) path has weight at most $n-1$. Using this improvement over Lemma 1 in Theorem 3, we get the following.

Lemma 2. *For an envy-freeable allocation* $\mathbf{A}$ *that is envy-free up to one good, no path in* $G_{\mathbf{A}}$ *has weight more than* $\min(n-1, m)$*. Hence,* $\mathrm{sub}(\mathbf{A}) \leq (n-1) \cdot \min(n-1, m)$.

With an envy-freeable EF1 allocation, the subsidy requirement becomes independent of the number of goods, at the expense of becoming quadratic in the number of agents. However, it is not even clear that an envy-freeable EF1 allocation always exists. For the special cases of binary and identical valuations, we show that it does, and in fact, picking a specific EF1 allocation that satisfies other properties allows achieving the optimal subsidy requirement of $n-1$.

Binary Valuations Recall that with binary valuations, we have $v_i(g) \in \{0, 1\}$ for all $i \in \mathcal{N}$ and $g \in \mathcal{M}$. We say that agent i *likes* good g if $v_i(g) = 1$. An allocation $\mathbf{A}$ is *non-wasteful* if each good is allocated to an agent who likes it. Note that because the valuations are binary, non-wasteful is equivalent to Pareto efficiency. For binary valuations, it is easy to see that every non-wasteful allocation is envy-freeable as it satisfies condition (b) of Theorem 1.

Algorithms such as the round-robin method and maximum Nash welfare (MNW) are known to produce non-wasteful EF1 allocations [11]. The round-robin method, given an agent ordering, allows agents to pick goods one-by-one according to the ordering in a cyclic fashion. The MNW algorithm finds the largest set of agents that can simultaneously receive positive utility and returns an allocation maximizing the product of their utilities.

Using a non-wasteful EF1 allocation, we can reduce the $O(mn)$ subsidy requirement to $O(n^2)$. This is the best we can do using the round-robin method with an arbitrary agent ordering (see the example in Appendix B.3). However, we show that the non-wasteful EF1 allocation returned by the MNW algorithm is special as it requires a subsidy of at most $n-1$, meeting the lower bound from Theorem 4.

Theorem 5. *For binary valuations, an allocation produced by the maximum Nash welfare algorithm is envy-freeable and requires at most $n-1$ subsidy.*

Proof. Let $\mathbf{A}$ be an allocation returned by the MNW algorithm. It is easy to see that $\mathbf{A}$ is non-wasteful, and hence, envy-freeable. Next, we show that any path in $G_{\mathbf{A}}$ has weight at most 1. This implies a subsidy requirement of at most $n-1$ using the same argument as in the proof of Theorem 3.

First, without loss of generality, we assume that each good is liked by at least one agent; if there are goods that are not liked by any agent, we could disregard them in the steps below and allocate them arbitrarily. We already argued that the non-wasteful allocation produced by the MNW algorithm is envy-freeable. Since it assigns each good to an agent who likes it, we have $v_i(A_i) = |A_i|$ for all $i \in \mathcal{N}$ and $v_i(A_j) \leq |A_j|$ for all $i, j \in \mathcal{N}$. It follows that $w(i,j) = v_i(A_j) - v_i(A_i) \leq |A_j| - |A_i|$ for all $i, j \in \mathcal{N}$.

Suppose for a contradiction that there exists a path P^* in $G_{\mathbf{A}}$ such that $w(P^*) > 1$. Because weights are integral, this implies $w(P^*) \geq 2$. Now, we make the following claim; the proof is given in Appendix C.

Claim. There exists a subpath P of P^* with no negative-weight edges and $w(P) \geq 2$.

Without loss of generality, we further assume that the first edge of P has a positive weight (otherwise we could consider the subpath of P starting at its first positive-weight edge). Let $P = (i_1, \ldots, i_k)$. We want to prove two claims: (a) $|A_{i_k}| \geq |A_{i_1}| + 2$, and (b) for each $t \in [k-1]$, there exists a good $g \in A_{i_{t+1}}$ which agent i_t likes.

For claim (a), recall that for each $t \in [k-1]$, we have $w(i_t, i_{t+1}) \leq |A_{i_{t+1}}| - |A_{i_t}|$. Summing over $t \in [k-1]$, we get that $|A_{i_k}| - |A_{i_1}| \geq w(P) \geq 2$, as desired.

Claim (b) holds for $t = 1$ because the first edge has weight $w(i_1, i_2) = v_{i_1}(A_{i_2}) - v_{i_1}(A_{i_1}) > 0$, implying $v_{i_1}(A_{i_2}) > 0$. For $t \in \{2, \ldots, k-1\}$, using the argument above, we have $|A_{i_t}| - |A_{i_1}| \geq w(i_1, \ldots, i_t) \geq w(i_1, i_2) > 0$. Hence, $v_{i_t}(A_{i_t}) = |A_{i_t}| > 0$. This, along with $w(i_t, i_{t+1}) = v_{i_t}(A_{i_{t+1}}) - v_{i_t}(A_{i_t}) \geq 0$, implies $v_{i_t}(A_{i_{t+1}}) > 0$.

Given the two claims, we derive a contradiction to the fact that $\mathbf{A}$ is returned by the MNW algorithm. Suppose we take a good from $A_{i_{t+1}}$ that agent i_t likes — it exists due to claim (b) — and add it to A_{i_t} for each $t \in [k-1]$. In the resulting allocation, the utility to agent i_k decreases by 1, the utility to agent i_1 increases by 1, and the utility to every other agent remains constant. Since agent i_k had at least 2 more utility than agent i_1 due to claim (a), it is easy to see that the resulting allocation would either give positive utility to strictly more agents (if i_1 had zero utility in the beginning) or strictly increase the product of utilities to the agents with positive utility. Both of these contradict the fact that $\mathbf{A}$ was returned by the MNW algorithm. Hence, every path in $G_{\mathbf{A}}$ has weight at most 1, which implies the desired result. □

Note that in this proof, along with non-wastefulness, the only property of the MNW algorithm that we used was the following: given allocations $\mathbf{A}^1$ and $\mathbf{A}^2$ such that for some agents $i, j \in \mathcal{N}$, $v_i(A_i^1) \geq v_j(A_j^1)+2$, $v_i(A_i^2) = v_i(A_i^1)-1$, $v_j(A_j^2) = v_j(A_j^1)+1$, and $v_k(A_k^1) = v_k(A_k^2)$ for all $k \in \mathcal{N} \setminus \{i, j\}$, the algorithm cannot return $\mathbf{A}^1$. This property as well as non-wastefulness are implied by the Pigou-Dalton principle [23]. Hence, the result holds for every algorithm which satisfies this principle, including the leximin rule.[2]

This proof leverages several ideas from the literature. Claim (a) shares similarities with a property of MNW allocations established by Darmann and Schauer [12], while the trick of passing goods along a path using claim (b) was also used by Barman et al. [3] to show that an MNW allocation can be computed efficiently for binary valuations. Thus, for binary valuations, we can efficiently compute an allocation which needs at most $n-1$ subsidy.

While the MNW algorithm achieves the optimal worst-case subsidy bound, it does not minimize the subsidy required on every instance. It is easy to construct instances where envy-free allocations exist but the MNW algorithm produces an allocation which requires as much as $n-2$ subsidy (see the example in Appendix B.3).

What is the complexity of computing the minimum subsidy required in a given allocation problem? As argued before, we can reduce the problem of checking the existence of an envy-free allocation to the problem of computing the minimum subsidy required. It is not difficult to see that the converse holds too. We can compute the minimum subsidy required by adding a unit subsidy at a time, and checking the existence of an envy-free allocation. The proof of the next result is given in Appendix C.

Proposition 2. *For binary valuations, the problems of computing the minimum subsidy required and checking the existence of an envy-free allocation are Turing-equivalent.*

Unfortunately, to the best of our knowledge, it is an open question whether existence of an envy-free allocation can be checked efficiently for binary valuations. However, the complexity of a closely related problem is known. Bouveret and Lang [9] show that checking the existence of a non-wasteful envy-free allocation with binary valuations is an NP-complete problem. Using the same argument as before, we have the following.

Corollary 1. *For binary valuations, it is NP-hard to compute the minimum subsidy required to achieve envy-freeness using a non-wasteful allocation.*

Identical Valuations With identical valuations, we denote the common valuation function of the agents by v. In this case, the utilitarian welfare $\sum_{i \in \mathcal{N}} v(A_i) = v(\mathcal{M})$ is constant. This implies that every allocation is Pareto efficient. Hence, by condition (b) of Theorem 1, every allocation $\mathbf{A}$ is envy-freeable.

[2] The leximin rule finds an allocation that maximizes the minimum utility, subject to that maximizes the second minimum utility, and so on.

Given an allocation $\mathbf{A}$, the optimal payment vector is given by $p_i^*(\mathbf{A}) = \max_{j\in\mathcal{N}} v(A_j) - v(A_i)$ for all $i \in \mathcal{N}$. To see this, note that each agent i requires payment at least $p_i^*(\mathbf{A})$ to not envy the agent with the highest utility. Conversely, $(\mathbf{A}, p_i^*(\mathbf{A}))$ is envy-free as every agent has the same value for all agents' allocations. Thus, $\text{sub}(\mathbf{A}) = n \cdot \max_{j\in\mathcal{N}} v(A_j) - v(\mathcal{M})$. Therefore, minimizing subsidy is equivalent to minimizing the maximum value of any bundle, which is the well-known NP-complete multiprocessor scheduling problem.

Proposition 3. *With identical valuations, every allocation is envy-freeable. An allocation minimizes the subsidy required if and only if it minimizes the maximum utility to any agent. Computing such an allocation is an NP-hard problem.*

What if we simply wanted to achieve the optimal worst-case upper bound of $n-1$ instead of minimizing the subsidy on every instance? For binary valuations, we achieved this by efficiently choosing a *specific* envy-freeable EF1 allocation — namely, the one produced by the MNW algorithm. For identical valuations, it is easy to see that *any* envy-freeable EF1 allocation $\mathbf{A}$ suffices as $p_i^*(\mathbf{A}) = \max_{j\in\mathcal{N}} v(A_j) - v(A_i)$ is at most 1 for each $i \in \mathcal{N}$ and is zero for some agent. Since we can compute an EF1 allocation efficiently, we have the following.

Proposition 4. *With identical valuations, we can efficiently compute an allocation which requires at most $n-1$ subsidy.*

Returning to General Valuations Recall that in the worst case, we need at least $n-1$ subsidy (Theorem 4). For the special cases of binary and identical valuations, we achieved this optimal bound by finding a special envy-freeable and EF1 allocation, respectively. For general valuations, the problem is that it is not clear if an envy-freeable EF1 allocation is even guaranteed to exist.

Most of the algorithms known in the literature that achieve EF1 are scale-free [11, 20], that is, multiplying an agent's valuation by a scalar does not affect the allocation returned. It is easy to see that such algorithms cannot always return an envy-freeable allocation.

Of these algorithms, the round-robin method is of special interest. With a fixed agent ordering, it is scale-free. But what if we chose the *right* agent ordering in a non-scale-free way? We show that this indeed works for two agents. The proof of the next result is provided in Appendix C.

Theorem 6. *When $n = 2$, there exists an agent ordering such that the allocation returned by the round-robin method with that ordering is envy-freeable and requires at most 1 subsidy.*

Note that this achieves the optimal bound of $n-1$ for $n = 2$ agents. Unfortunately, this method does not work for $n \geq 3$ agents. In our counterexample (see Appendix B.3), while the round-robin method fails to produce an envy-freeable EF1 allocation with any agent ordering, there still exists an envy-freeable EF1 allocation. This leads us to the following conjecture.

Conjecture 1. There always exists an envy-freeable allocation that is envy-free up to one good.

If this conjecture is true, then by Lemma 2, we know that the minimum subsidy required in the worst case is $O(n^2)$ (thus independent of m). We conjecture further that the lower bound of $n-1$ can be achieved.

Conjecture 2. There always exists an envy-freeable allocation which requires at most $n-1$ subsidy.

In fact, it may be possible that a subsidy of at most $n-1$ can always be achieved through an envy-freeable EF1 allocation.

5 Experiments

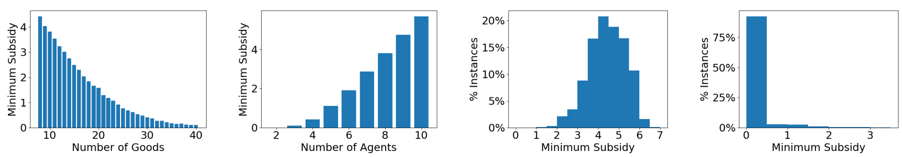

(a) Avg,n=8 (b) Avg,m=10 (c) Dist,n=8,m=8 (d) Dist,n=8,m=40

Fig. 1. The minimum subsidy required in our simulations. Figures (a) and (b) show the minimum subsidy averaged across instances as functions of m and n, respectively. Figures (c) and (d) show the distribution of minimum subsidy for fixed n and m.

In this section, we empirically study the minimum subsidy required in the average case. We compute the minimum subsidy required across all allocations by solving an integer linear program using CPLEX.

To generate synthetic data, we consider instances with $2 \leq n \leq 8$ and $n \leq m \leq 5n$. For each (n, m), we sample $1,000$ instances as follows: For each good g we sample $v^*(g)$ from an exponential distribution with mean 30 and $\sigma^*(g)$ from an exponential distribution with mean 5. Then, for each agent i and good g, we draw $v_i(g)$ from a truncated normal distribution, which has mean $v^*(g)$ and standard deviation $\sigma^*(g)$, and is truncated below at 0.

In addition, we obtained $3,535$ real-world fair division instances from a popular fair division website Spliddit.org. These instances have divisible as well as indivisible goods, from 2 to 15 agents, and from 2 to 96 goods. While Spliddit data does not match our model as agents are forced to report valuations that sum to a constant, we believe that it still provides a valuable empirical perspective.

We begin by noting that none of the $114,000$ synthetic instances or $3,535$ real-world instances required a subsidy of more than $n-1$, which is evidence in support of Conjecture 2.

In our synthetic experiments, we see that fixing the number of agents, the minimum subsidy required reduces on average as the number of goods increases

(Figure 1(a)). On the other hand, fixing the number of goods, the minimum subsidy required (almost linearly) increases on average as the number of agents increases (Figure 1(b)). These results are in part due to the fact that the probability of existence of an envy-free allocation (i.e., of requiring no subsidy) increases with more goods but decreases with more agents [14]. Next, we dive into the distribution of the minimum subsidy required, presented in Figure 1(c) for $n = m = 8$ and in Figure 1(d) for $n = 8$ and $m = 40$. Again, with more goods, the distribution quickly skews towards requiring little to no subsidy.

Finally, on the real-world data obtained from Spliddit, 68% of the instances required no subsidy (i.e., admitted envy-free allocations), while 93% of the instances required a subsidy of at most 1. Thus, in practice, the amount of subsidy needed to eliminate envy is most likely no greater than the maximum value that any agent places on a single good.

6 Discussion

We have examined the minimum subsidy required both in cases when an allocation is given to us and when it can be chosen. In the former case, we have shown how to compute the minimum subsidy exactly; in both cases, we have provided several useful bounds for cases of interest. However, a number of directions remain open for further research. Perhaps the most immediate question is to settle our two conjectures from Section 4.2. Specifically, it may be possible to adapt the iterative algorithm of Lipton et al. [20] to select the good to be allocated in each iteration in a non-scale-free way and achieve the optimal bound of $n - 1$ subsidy. Settling the complexity of checking the existence of an envy-free allocation for binary valuations is also an important open question. Finally, it would be interesting to extend this framework to non-additive valuations.

More broadly, while we modeled the divisible good as external subsidy throughout the paper, our results also have implications for other models of introducing monetary payments. For example, when no subsidy is available but monetary transfers among agents are possible, we would like to find *budget-balanced transfers*, $\mathbf{p}$ where $\sum_{i \in \mathcal{N}} p_i = 0$. It is easy to show that computing the optimal payment vector from Theorem 2 and then reducing the payment to each agent by the average payment finds budget-balanced transfers which minimize the maximum amount that any agent has to pay. Alternatively, one could consider a model where each agent pays to receive goods ($p_i \leq 0$ for each i). It is again easy to show that we can efficiently minimize the total payment collected in a manner similar to Theorem 2. It would be interesting to study other natural objective functions (e.g., minimizing the number of agents that have a non-zero payment) in such models.

Bibliography

[1] Alkan, A., Demange, G., Gale, D.: Fair allocation of indivisible goods and criteria of justice. Econometrica **59**(4), 1023—1039 (1991)

[2] Barman, S., Krishnamurthy, S.K., Vaish, R.: Finding fair and efficient allocations. In: Proceedings of the 19th ACM Conference on Economics and Computation (EC), pp. 557–574 (2018)

[3] Barman, S., Krishnamurthy, S.K., Vaish, R.: Greedy algorithms for maximizing Nash social welfare. In: Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), pp. 7–13 (2018)

[4] Berliant, M., Dunz, K., Thomson, W.: On the fair division of a heterogeneous commodity. Journal of Mathematical Economics **21**, 201–216 (1992)

[5] Beviá, C., Quinzii, M., Silva, J.A.: Buying several indivisible goods. Mathematical Social Sciences **37**(1), 1–23 (1999)

[6] Bikhchandani, S., Mamer, J.W.: Competitive equilibrium in an exchange economy with indivisibilities. Journal of Economic Theory **74**(2), 385–413 (1997)

[7] Bilo, V., Ioannis, C., Flammini, M., Ayumi, I., Monaco, G., Dominik, P., Vinci, C., Zwicker, W.S.: Almost envy-free allocations with connected bundles. In: Proceedings of the 10th Innovations in Theoretical Computer Science Conference (ITCS), pp. 1–21, 14 (2019)

[8] Bouveret, S., Cechlárová, K., Elkind, E., Igarashi, A., Peters, D.: Fair division of a graph. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI), pp. 135–141 (2017)

[9] Bouveret, S., Lang, J.: Efficiency and envy-freeness in fair division of indivisible goods: logical representation and complexity. Journal of Artificial Intelligence Research **32**, 525–564 (2008)

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

[11] Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J.: The unreasonable fairness of maximum Nash welfare. In: Proceedings of the 17th ACM Conference on Economics and Computation (EC), pp. 305–322 (2016)

[12] Darmann, A., Schauer, J.: Maximizing nash product social welfare in allocating indivisible goods. European Journal of Operational Research **247**(2), 548–559 (2015)

[13] Demange, G., Gale, D.: The strategy structure of two-sided matching markets. Econometrica pp. 873–888 (1985)

[14] Dickerson, J.P., Goldman, J., Karp, J., Procaccia, A.D., Sandholm, T.: The computational rise and fall of fairness. In: Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI), pp. 1405–1411 (2014)

[15] Edmonds, J., Karp, R.M.: Theoretical improvements in algorithmic efficiency for network flow problems. Journal of the ACM (JACM) **19**(2), 248–264 (1972)
[16] Eisenberg, E., Gale, D.: Consensus of subjective probabilities: The pari-mutuel method. The Annals of Mathematical Statistics **30**(1), 165–168 (1959)
[17] Foley, D.: Resource allocation and the public sector. Yale Economics Essays **7**, 45–98 (1967)
[18] Haake, C.J., Raith, M.G., Su, F.E.: Bidding for envy-freeness: A procedural approach to n-player fair-division problems. Social Choice and Welfare **19**(4), 723–749 (2002)
[19] Klijn, F.: An algorithm for envy-free allocations in an economy with indivisible objects and money. Social Choice and Welfare **17**(2), 201–215 (2000)
[20] Lipton, R.J., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of indivisible goods. In: Proceedings of the 6th ACM Conference on Economics and Computation (EC), pp. 125–131 (2004)
[21] Maskin, E.S.: On the fair allocation of indivisible goods. In: Arrow and the Foundations of the Theory of Economic Policy, pp. 341–349 (1987)
[22] Meertens, M., Potters, J., Reijnierse, H.: Envy-free and pareto efficient allocations in economies with indivisible goods and money. Mathematical Social Sciences **44**(3), 223–233 (2002)
[23] Moulin, H.: Fair Division and Collective Welfare. MIT Press (2004)
[24] Ohseto, S.: Characterizations of strategy-proof and fair mechanisms for allocating indivisible goods. Economic Theory **29**(1), 111–121 (2006)
[25] Pazner, E., Schmeidler, D.: Egalitarian equivalent allocations: A new concept of economic equity. Quarterly Journal of Economics **92**(4), 671–687 (1978)
[26] Plaut, B., Rougligarden, T.: Almost envy-freeness with general valuations. In: Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 2584–2603 (2018)
[27] Procaccia, A.D., Wang, J.: Fair enough: Guaranteeing approximate maximin shares. In: Proceedings of the 14th ACM Conference on Economics and Computation (EC), pp. 675–692 (2014)
[28] Quinzii, M.: Core and competitive equilibria with indivisibilities. International Journal of Game Theory **13**(1), 41–60 (1984)
[29] Steinhaus, H.: The problem of fair division. Econometrica **16**, 101–104 (1948)
[30] Su, F.E.: Rental harmony: Sperner's lemma in fair division. American Mathematical Monthly **106**(10), 930–942 (1999)
[31] Svensson, L.G.: Large indivisibles: An analysis with respect to price equilibrium and fairness. Econometrica **51**(4), 939–954 (1983)
[32] Varian, H.: Equity, envy and efficiency. Journal of Economic Theory **9**, 63–91 (1974)
[33] Weller, D.: Fair division of a measurable space. Journal of Mathematical Economics **14**(1), 5–17 (1985)

Appendix

A Lattice Structure of Envy-Eliminating Payments

We argue that for an envy-freeable allocation $\mathbf{A}$, the set of envy-eliminating payment vectors $\mathcal{P}(\mathbf{A})$ has a lattice structure.

In this lattice, two payment vectors $\mathbf{x}$ and $\mathbf{y}$ are compared using the partial order relation $\preceq$, where $\mathbf{x} \preceq \mathbf{y}$ if $x_i \leq y_i$ for every agent $i \in \mathcal{N}$. We need to show that for every two payment vectors $\mathbf{x}, \mathbf{y} \in \mathcal{P}(\mathbf{A})$, there exist a greatest lower bound and a least upper bound.

We argue that the greatest lower bound is given by ℓ such that $\ell_i = \min(x_i, y_i)$ for each $i \in \mathcal{N}$ and the least upper bound is given by $\mathbf{u}$ such that $u_i = \max(x_i, y_i)$ for each $i \in \mathcal{N}$. We argue that ℓ is the greatest lower bound; the argument for the least upper bound is similar.

Note that $\ell \preceq \mathbf{x}$ and $\ell \preceq \mathbf{y}$ holds as $\ell_i = \min(x_i, y_i) \leq x_i$ (resp. y_i). Additionally, it holds that for any $\mathbf{p} \in \mathcal{P}(\mathbf{A})$ such that $\mathbf{p} \preceq \mathbf{x}$ and $\mathbf{p} \preceq \mathbf{y}$, we have $\mathbf{p} \preceq \ell$ since $p_i \leq x_i$ and $p_i \leq y_i$ implies $p_i \leq \min(x_i, y_i) = \ell_i$. Hence, to complete the argument, we simply need to show that $\ell \in \mathcal{P}(\mathbf{A})$. Consider agents $i, j \in \mathcal{N}$. Without loss of generality, assume that $x_i \leq y_i$. It follows that $v_i(A_i) + \ell_i = v_i(A_i) + x_i \geq v_i(A_j) + x_j \geq v_i(A_j) + \ell_j$, where the first inequality holds because $\mathbf{x} \in \mathcal{P}(\mathbf{A})$. Hence, $\ell \in \mathcal{P}(\mathbf{A})$. It is also easy to show that $\mathbf{p}^*(\mathbf{A})$ is the unique lowest element of this lattice.

B Missing Details

B.1 Minimizing Subsidy for a Given Allocation via an LP

$$
\begin{aligned}
&\text{Minimize } \textstyle\sum_{i \in \mathcal{N}} p_i \\
&\text{s.t.} \\
&v_i(A_i) + p_i \geq v_i(A_j) + p_j, \forall i, j \in \mathcal{N} \\
&p_i \geq 0, \forall i \in \mathcal{N}
\end{aligned}
\tag{1}
$$

B.2 Minimizing Subsidy Across Allocations via an ILP

$$
\begin{aligned}
&\text{Minimize } \textstyle\sum_{i \in \mathcal{N}} p_i \\
&\text{s.t.} \\
&\textstyle\sum_{g \in \mathcal{M}} v_i(g) x_{i,g} + p_i \geq \sum_{g \in \mathcal{M}} v_i(g) x_{j,g} + p_j, \forall i, j \in \mathcal{N} \\
&\textstyle\sum_{i \in \mathcal{N}} x_{i,g} = 1, \forall g \in \mathcal{M} \\
&x_{i,g} \in \{0, 1\}, \forall i \in \mathcal{N}, g \in \mathcal{M} \\
&p_i \geq 0, \forall i \in \mathcal{N}
\end{aligned}
\tag{2}
$$

B.3 Missing Examples

Example 1 (Round-Robin Requires $\Omega(n^2)$ Subsidy for Binary Valuations). Consider the following allocation problem with n agents having binary valuations over $m = n(n+1)/2$ goods denoted $g_1, \ldots, g_m$.

Let us partition the goods as follows: for $i \in [n]$, let

$$X_i = \left\{ g_t : \frac{i(i-1)}{2} + 1 \leq t \leq \frac{i(i+1)}{2} \right\}.$$

Note that $|X_i| = i$ for each $i \in [n]$. Let $X_{n+1} = \emptyset$. The valuations are as follows: each agent $i \in \mathcal{N}$ likes each good in $X_i \cup X_{i+1}$ and does not like every other good.

Suppose that we run the round-robin with agent ordering $n, \ldots, 1$. Further, suppose that each agent i, in each turn that she gets, picks a good from X_i. Then, it is easy to see that the allocation produced will assign to each agent $i \in \mathcal{N}$ the set of goods X_i. Under this allocation, agent i will envy agent $i+1$ by one good for each $i \in [n-1]$. Thus, the maximum weight of a path starting at agent i will be $n - i$, and the amount of subsidy required will be $\sum_{i=1}^{n}(n-i) = \Omega(n^2)$.

Example 2 (MNW Does Not Minimize Subsidy Instance-Wise for Binary Valuations). Consider an allocation problem with n agents and $n+2$ goods. The first $n-1$ agents value the first n goods at 1, and the last two goods at 0. Agent n values every good at 1.

In this case, the MNW algorithm returns an allocation in which agent n receives the last two goods, one of the first $n-1$ agents receives two of the first n goods, and each remaining agent receives a single good. It is easy to see that this requires $n-2$ subsidy to eliminate envy among the first $n-1$ agents.

In contrast, giving the last three goods to agent n and giving each remaining agent a single good would result in an envy-free allocation, requiring zero subsidy.

Example 3 (Round-Robin Does Not Produce Envy-Freeable EF1 Allocation). Consider the following instance with 3 agents and 4 goods, where the valuations are given by the following matrix.

$$\begin{pmatrix} 5 & 2 & 3 & 32 \\ 23 & 1 & 7 & 38 \\ 15 & 2 & 1 & 23 \end{pmatrix}$$

One can check that regardless of the agent ordering, the round-robin method produces an allocation that is not envy-freeable. However, in this example, there exists an envy-freeable EF1 allocation (for instance, consider the allocation where agent 1 receives g_2 and g_3, agent 2 receives g_1, and agent 3 receives g_4).

C Missing Proofs

Proof (Proposition 1). Given $\mathbf{A}$, create a weighted complete bipartite graph G with n nodes on each side. Set the weight of the edge from node i on the left to node j on the right as $v_i(A_j)$. This graph can be constructed in $O(mn + n^2)$ time. Then, we can compute its maximum weight bipartite matching in $O(n^3)$ time [15]. Condition (b) from Theorem 1 implies that $\mathbf{A}$ is envy-freeable if and only if the utilitarian welfare under $\mathbf{A}$, which can be computed in $O(nm)$ time, is at least the weight of this matching. Hence, the total running time is $O(mn+n^3)$.

Alternatively, we can use condition (c) from Theorem 1, and check the existence of positive-weight cycles in $G_{\mathbf{A}}$. This can also be accomplished in $O(nm + n^3)$ time by running the Floyd-Warshall (which can check the existence of negative-weight cycles) on the graph obtained by negating all edge weights in $G_{\mathbf{A}}$. The graph construction takes $O(nm)$ time and the Floyd-Marshall algorithm runs in $O(n^3)$ time. □

Proof (Lemma 2). As in Lemma 1, since $G_{\mathbf{A}}$ has no positive-weight cycles, we only need to worry about simple paths on which no agent appears twice. Consider any such path $(i_1, \dots, i_k)$. From Lemma 1, we already know that its weight is at most m.

The fact that $\mathbf{A}$ is EF1 implies that each edge in $G_{\mathbf{A}}$ has weight at most 1. Since a simple path can have at most $n - 1$ edges, its weight is at most $n - 1$. Using the same argument as in the proof of Theorem 3, we get that the subsidy required is at most $(n - 1) \cdot \min(n - 1, m)$. □

Proof (Claim from Theorem 5). Let $P^* = (i_1^*, \dots, i_r^*)$. Let us decompose P^* into consecutive segments $P_1, \dots, P_t$ that alternate between having either only negative or only non-negative edge weights. More formally, we break P^* at every node i_t^* such that one of $w(i_{t-1}^*, i_t^*)$ and $w(i_t^*, i_{t+1}^*)$ is negative while the other is non-negative. Since weights are integral, each segment consisting of only negative-weight edges has weight at most -1. If each segment consisting of only non-negative-weight edges has weight at most 1, then it is easy to see that we would have $w(P^*) \leq 1$, which would be a contradiction. Hence, there must exist a segment P consisting of only non-negative-weight edges with $w(P) \geq 2$. □

Proof (Proposition 2). Note that in an allocation problem with binary valuations, the minimum subsidy required is integral. This is because for every allocation $\mathbf{A}$, edges in $G_{\mathbf{A}}$ have integral weight. Hence, from Theorem 2, $\text{sub}(\mathbf{A})$ is integral. Thus, the minimum subsidy required, i.e., $\min_{\mathbf{A}} \text{sub}(\mathbf{A})$ is also integral. Further, when the minimum subsidy is allocated to achieve envy-freeness, each agent must receive integral payment (Theorem 2). Hence, we can treat a subsidy of k as k indivisible goods which all agents value at 1.

Now, we can compute the minimum subsidy required as follows: we can add, one-by-one, indivisible goods that all agents value at 1, and check for the existence of envy-free allocations. The least number of goods we need to add for envy-free allocations to exist is the minimum subsidy required. Theorem 5 implies that we need to add at most $n - 1$ goods. Hence, we can compute the minimum subsidy required by making n calls to an algorithm for checking the existence of an envy-free allocation. We have already established the reduction in the other direction. □

Proof (Theorem 6). For $i \in \{1, 2\}$, let v_i^k denote the k^{th} largest value that agent i has for any good, $v_i^{odd} = \sum_{odd k} v_i^k$, and $v_i^{even} = \sum_{even k} v_i^k$. Note that $v_i^{odd} + v_i^{even} = v_i(\mathcal{M})$.

Let $\mathbf{A}^{12}$ (resp., $\mathbf{A}^{21}$) denote the allocation produced by the round-robin method when agent 1 (resp., agent 2) goes first. Note that $v_1(A_1^{12}) \geq v_1^{odd}$. This is because when agent 1 chooses a good in round $2k+1$, one of her $2k+1$ most valuable goods must be available. Hence, in this round, she must pick a good for which her value is at least v_1^{2k+1}. Summing over all k, we get that $v_1(A_1^{12}) \geq v_1^{odd}$. Similarly, $v_2(A_2^{12}) \geq v_2^{even}$. Hence, we also have $v_1(A_2^{12}) \leq v_1(\mathcal{M}) - v_1^{odd} = v_1^{even}$ and $v_2(A_1^{12}) \leq v_2(\mathcal{M}) - v_2^{even} = v_2^{odd}$. Now, if $\mathbf{A}^{12}$ is not envy-freeable, then by Theorem 1, we have

$$v_1(A_2^{12}) + v_2(A_1^{12}) > v_1(A_2^{12}) + v_2(A_1^{12})$$
$$\Rightarrow v_1^{even} + v_2^{odd} > v_1^{odd} + v_2^{even}. \tag{3}$$

Similarly, if $\mathbf{A}^{21}$ is not envy-freeable, then we have

$$v_1^{odd} + v_2^{even} > v_1^{even} + v_2^{odd}. \tag{4}$$

However, Equations (3) and (4) contradict each other. Hence, at least one of $\mathbf{A}^{12}$ and $\mathbf{A}^{21}$ must be envy-freeable.

Without loss of generality, suppose $\mathbf{A}^{12}$ is envy-freeable. Note that agent 1 does not envy agent 2, and agent 2 envies agent 1 by at most one good (because round-robin satisfies EF1). Hence, eliminating envy requires a total subsidy of at most 1. □