ECAI 2024
U. Endriss et al. (Eds.)
© 2024 The Authors.
This article is published online with Open Access by IOS Press and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0).
doi:10.3233/FAIA240920

The Degree of Fairness in Efficient House Allocation

Hadi Hosseini[a,*], **Medha Kumar**[a,**] **and Sanjukta Roy**[b,***]

[a]Pennsylvania State University, USA
[b]University of Leeds, UK

Abstract. The classic house allocation problem is primarily concerned with finding a matching between a set of agents and a set of houses that guarantees some notion of economic efficiency (e.g. utilitarian welfare). While recent works have shifted focus on achieving fairness (e.g. minimizing the number of envious agents), they often come with notable costs on efficiency notions such as utilitarian or egalitarian welfare. We investigate the trade-offs between these welfare measures and several natural fairness measures that rely on the number of envious agents, the total (aggregate) envy of all agents, and maximum total envy of an agent. In particular, by focusing on envy-free allocations, we first show that, should one exist, finding an envy-free allocation with maximum utilitarian or egalitarian welfare is computationally tractable. We highlight a rather stark contrast between utilitarian and egalitarian welfare by showing that finding utilitarian welfare maximizing allocations that minimize the aforementioned fairness measures can be done in polynomial time while their egalitarian counterparts remain intractable (for the most part) even under binary valuations. We complement our theoretical findings by giving insights into the relationship between the different fairness measures and by conducting empirical analysis.

1 Introduction

The classic house allocation problem is primarily concerned with assigning a set of houses (or resources) to a set of agents based on their preferences over houses such that each agent receives at most one house. It was motivated by a variety of applications such as kidney exchange [23] or labour market [14] where agents are initially endowed with houses or houses have to be distributed afresh among agents.[1] While this model was primarily studied for designing incentive compatible mechanisms [26, 1] along with some notion of economic efficiency, recent works have shifted focus to the issues of fairness such as envy-freeness (EF), which requires that every agent weakly prefers its allocated house to that of every other agent. An envy-free allocation may not always exist: for example, consider two agents who both like the same house. The fair division literature contains a variety of *approximate* envy measures (e.g. envy-free up to one item [18, 7]) that cannot appropriately be utilized here due to the 'one house per agent' constraint.

* Email: hadi@psu.edu
** Email: mkumar@psu.edu
*** Email: s.roy@leeds.ac.uk

[1] This problem is commonly known as "Shapley-Scarf Housing Markets" when agents are initially endowed with houses. The goal is often finding mutually-beneficial exchanges that lead to efficient stable allocations (see [24, 22]).

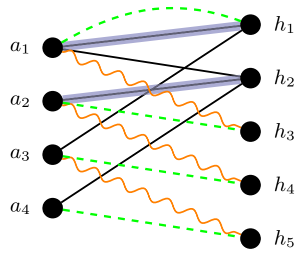

Figure 1: An envy-free allocation of maximum size is shown in wavy orange; a minimum #envy complete allocation is shown in dashed green; no envy-free allocation has the maximum welfare 2; and the allocation in blue achieves minimum #envy among the matchings that have the maximum welfare.

An orthogonal, but more suitable, approach is measuring the 'degree of envy' [8] among the agents by counting the number of envious agents or the total aggregate envy experienced. In this vein, recent works have investigated the existence of envy-free house allocations under ordinal preferences [9], maximum-size envy-free allocation [2], and complete allocations that minimizes the number of envious agents [16, 19] or those that minimize the total aggregate envy among agents [11]. Yet, these approaches often take a toll on efficiency notions such as size (the number of allocated agents), utilitarian welfare (the sum of agents' values), or egalitarian welfare (the value of the worst off agent). Let us illustrate some of these nuances with an example.

Example 1 (**Fairness in House Allocation**). *Consider an instance with four agents* $\{a_1, a_2, a_3, a_4\}$ *and five houses* $\{h_1, h_2, \ldots, h_5\}$ *and binary valuations. For the ease of exposition, we use a graphical representation of the problem as shown in Figure 1, where solid lines* $(a_1, h_1), (a_1, h_2), (a_2, h_2), (a_3, h_1)$*, and* (a_4, h_2) *indicate an agent has a positive valuation (value of* 1*) for the house.*

A maximum-size envy-free allocation (shown in wavy orange lines) assigns houses h_3*,* h_4*, and* h_5 *to agents* a_1*,* a_2*, and* a_3 *respectively. While the size of the maximum size envy-free allocation is three, its utilitarian welfare is zero. Both the maximum size envy-free allocation and the envy-free allocation of maximum utilitarian welfare, have zero utilitarian welfare. A maximum utilitarian welfare allocation has utilitarian welfare two. An allocation with welfare two that minimizes the number of envious agents (highlighted in blue) has two envious agents* a_3 *and* a_4*. There is no envy-free allocation with maximum utilitarian welfare, i.e., no envy-free allocation with utilitarian welfare two. A complete allocation (agent saturating) that minimizes the number of envious agents is shown by dashed green lines.*

In the above example, completeness refers to solutions that maximally utilize the resources, an efficiency property that is essential in many societal domains such as allocating public housing to families and refugee settlement [20, 5]. It turns out that finding a complete allocation that minimizes the number of envious agents (the

	max USW	max ESW	size $\geq k$
envy-free (EF)	P (Prop. 3)†	P (Thm. 6)	P (Prop. 3)⋆
min #envy	P (Thm. 2)	NP-c (Thm. 7)	$m > n$: NP-c§ $m \leq n$: P (Thm. 3)
min total envy	P (Thm. 4)	as hard as ($) (Thm. 8)	$m > n$: open ($) $m \leq n$: P (Cor. 1)‡
minimax total envy	open	NP-c (Thm. 9)	NP-c‡

Table 1: The summary of results for weighted instances with n agents, m houses, and $0 \leq k \leq \min\{m, n\}$. P and NP-c refer to polynomial time and NP-complete, respectively. ($) refers to the min total envy complete when $m > n$. Result(s) marked by § are due to [16], those marked by ‡ and † are shown by [19] and [2], respectively for binary valuations, and ⋆ is the result by [9] for $k = n$.

green allocation in the above example) is NP-complete even for binary valuations [16]. Moreover, a maximum-size envy-free allocation proposed by Aigner-Horev and Segal-Halevi [2] only considers allocations to positive valued houses and returns 'empty' otherwise, and the envy-free algorithm proposed by Gan et al. [9] returns 'none' because there is no complete envy-free allocation that assigns *every* agent to a house it likes. While these observations illustrate the intricate trade-offs between each approach, the relation between these notions and their computational aspects have remained unstudied.

1.1 Our Contributions

We study the computational problems of finding a fair house allocation *within the set of allocations that maximize an efficiency measure*. We focus on three well-studied notions of economic efficiency, namely, the size (i.e., the number of allocated agents), the utilitarian welfare, and the egalitarian welfare of the allocation. We study these notions under both binary and arbitrary positive valuations and investigate their interaction with various envy fairness measures. These include envy-freeness (EF), minimizing the number of envious agents (min #envy), and minimizing the total envy of all agents (min total envy), and minimizing the total envy of the most envious agent (minimax total envy). Table 1 summarizes our results.

Envy-free allocations. We first show that an envy-free allocation of maximum size can be computed efficiently under binary (Proposition 2) or arbitrary non-negative valuations (Theorem 1). By focusing on welfare-maximizing allocations, we show that, should one exist, finding an envy-free allocation that maximizes utilitarian welfare (USW) or egalitarian welfare (ESW) is computationally tractable (Proposition 3 and Theorem 6, resp.).

Utilitarian welfare. We show that finding an allocation with min #envy (Theorem 2), or min total envy (Theorem 4), is tractable under the added constraint of maximizing USW. Without the welfare constraint, the former problem has been proven to be NP-hard when we aim to find a complete allocation [16]. Additionally, we analyze the relationship between the size of an allocation and its USW. We obtain polynomial time algorithms for finding min #envy (Theorem 3), min total envy (Corollary 1) complete allocations when $m \leq n$.

Egalitarian welfare. To complement our study, we consider the well-established Rawlsian notion of egalitarian welfare. We first show that finding an envy-free allocation, when one exists, with maximum ESW can be done in polynomial time (Theorem 6). We highlight a contrast between egalitarian and utilitarian welfare by showing that when considering allocations that maximize the egalitarian welfare (as opposed to utilitarian ones), finding a min #envy allocation is NP-complete (Theorem 7). Moreover, it is NP-complete to find minimax total envy max ESW (Theorem 9). Finally, while the computational complexity of finding a complete allocation that minimizes the total envy remains open, we show an intriguing relation to its egalitarian welfare counterpart (Theorem 8) and conclude with complementary experimental observations under randomly generated valuations.

1.2 Related Work

Fairness in house allocations is a well studied problem. House allocations were first studied in social choice where the focus was on finding efficient and strategyproof allocations [22]. As the focus shifted towards fair allocations, increasingly algorithmic approaches were used. Achieving fairness by minimizing the number of envious agents or total envy has been studied previously by [9, 2, 16, 19]. Kamiyama [15] showed hardness of finding EF solutions for pairwise preferences. Belahcène et al. [4] studied a relaxed notion called ranked envy-freeness. The hardness and approximability of minimizing total envy was studied for housing allocation problems where agents are located on a graph [13, 11]. The egalitarian allocations have been studied under different names such as the classic makespan minimization problem in job scheduling [17], the Santa Claus problem [3], and in fair allocation of resources [18, 6]. In these settings the problem of maximizing the egalitarian welfare (worst-off agent) is shown to be NP-hard [6], giving rise to several approximation algorithms [25]. See the full version of our paper [12] for an extended discussion on related work. Our setting is crucially different from these works in two ways: in the house allocation problem each agent receives at most one house, and not all houses need to be assigned.

2 Model

An instance of the *house allocation problem* is represented by a tuple $\langle N, H, V \rangle$, where $N := \{1, 2, \dots, n\}$ is a set of n agents, $H := \{h_1, h_2, \dots, h_m\}$ is a set of m houses, and $V := (v_1, v_2, \dots, v_n)$ is a *valuation profile*. Each $v_i(h)$ indicates agent i's non-negative value for house $h \in H$. Thus, for $i \in N$ the value of a house $h \in H$ is $v_i(h) \geq 0$, and $v_i(\emptyset) = 0$. An instance is *binary* if for every $i \in N$ and every $h \in H$, $v_i(h) \in \{0, 1\}$; otherwise it is a *weighted* instance.

An allocation A is an injective mapping from agents in N to houses in H. For each agent $i \in N$, $A(i)$ is the house allocated to agent i given the allocation A and $v_i(A(i))$ is its value. Thus, for each $\{i, j\} \subseteq N$, $A(i) \cap A(j) = \emptyset$ and any house is allocated to at most one agent. The set of all such allocations is denoted by $\mathcal{A}$.[2]

Fairness. Given an allocation A, we say that agent i *envies* j if $v_i(A(j)) > v_i(A(i))$. The *amount (magnitude)* of this pairwise envy is captured by $\mathsf{envy}_{i,j}(A) := \max\{v_i(A(j)) - v_i(A(i)), 0\}$. Given an allocation A, the *total (aggregate) envy* of an agent i towards other agents is denoted by $\mathsf{envy}_i(A) = \sum_{j \in N} \mathsf{envy}_{i,j}(A)$.

An allocation A is *envy-free* (EF) if and only if for every pair of agents $i, j \in N$ we have $v_i(A(i)) \geq v_i(A(j))$, that is, $\mathsf{envy}_{i,j}(A) = 0$. Since envy-free allocations are not guaranteed to exist, we consider other plausible approximations to measure the 'degree of envy' [8].

[2] Note that in its graphical representation, this model differs from the classical bipartite matching problem as it allows for allocations along zero-valued edges (non-edges) in a graph.

Degrees of envy. An allocation is **min #envy** if it minimizes the number of envious agents, i.e. $\min_{A \in \mathcal{A}} \#\mathsf{envy}(A)$, where $\#\mathsf{envy}(A)$ is the number of envious agents, i.e., the size of the set $\{i \in N : \mathsf{envy}_{i,j}(A) > 0, \text{for some } j \in N\}$.

An allocation is **min total envy** if it minimizes the total envy of all agents, i.e. $\min_{A \in \mathcal{A}} \mathsf{total\ envy}(A)$, where $\mathsf{total\ envy}(A)$ is the total amount of envy of allocation A experienced by all agents, i.e., $\mathsf{total\ envy}(A) := \sum_{i \in N} \mathsf{envy}_i(A)$.

An allocation is **minimax total envy** if it minimizes the maximum aggregate amount of envy experienced by an agent, i.e. $\min_{A \in \mathcal{A}} \max_{i \in N} \mathsf{envy}_i(A)$.

In the above measures, an allocation is selected from the set of all feasible allocations $\mathcal{A}$. In the next section we discuss the reason behind restricting the set $\mathcal{A}$ to subsets that satisfy some measures of economic efficiency.

Social Welfare. Without any measures of social welfare, any empty allocation is vacuously envy-free, and consequently satisfies all four measures of fairness. Hence, we consider three notions of social welfare that measure the economic efficiency of allocations based on their size (number of assigned agents), utilitarian, or egalitarian welfare.

The **size** $|A|$ of an allocation, A, is simply the number of agents that are assigned to a house.[3] An allocation is **complete** if it either assigns a house to every agent (N-saturating) when $m \geq n$, or assigns every house to an agent (H-saturating) when $m < n$. Note that completeness is a weak efficiency requirement that does not take agents' valuations into account.

The **utilitarian welfare** of an allocation A is the sum of the values of individual agents, i.e. $\mathsf{USW}(A) := \sum_{i \in N} v_i(A)$. A maximum utilitarian welfare allocation is the one that maximizes the sum of the values, and can be found efficiently by computing a maximum-weight bipartite matching in the induced graph (a bipartite graph on $N \cup H$ where edge weights are given by the valuations V).

The **egalitarian welfare** of an allocation A is the value of the worst off agent among all agents in N, that is, $\mathsf{ESW}(A) := \min_{i \in N} v_i(A)$. A k-egalitarian welfare is the value of the worst off agent in a subset $S \subseteq N$ of agents of size $k = |S|$ such that $\mathsf{ESW}(A) := \min_{i \in S} v_i(A)$.

If it is possible to achieve a positive (non-zero) egalitarian welfare for all agents, any allocation that maximizes the egalitarian welfare (there could be multiple) can be selected. In the special case where every feasible N-saturating allocation (allocations of size n) has an egalitarian welfare of zero, we look for the largest subset of agents $S \subseteq N$ that can simultaneously receive a positive value, and select an allocation that maximizes the k-egalitarian welfare among these agents. In Example 1, a maximum egalitarian welfare allocation has welfare one and size two since any larger subset of agents will result in a egalitarian welfare of zero.

Computational Problems. Our main objective is to investigate the four fairness measures and various notions of welfare (i.e. economic efficiency). Thus, for each fairness-welfare pair, we define computational problems in the following way:

Given an instance of the house allocation problem, $\langle N, H, V \rangle$, find an allocation A that minimizes unfairness as measured by F within the set of all allocations that maximize an efficiency measure of E, where F is either min #envy, min total envy, or minimax total envy and E is either max size, max USW, or max ESW. The optimal allocation is termed as a *min F max E* allocation.

Some standard graph theoretic notations and algorithms that we use are provided in [12] for reference.

[3] We intentionally use the term 'size' instead of 'cardinality' to avoid confusion with matchings that *only* allow selection of positively valued (aka. 'liked') edges of a graph.

3 Maximum Size Envy-free Allocations

We start by considering envy-freeness as our main constraint. The goal is to find an envy-free allocation under various welfare measures. Clearly, an empty allocation is always an envy-free allocation of size zero. However, a non-empty envy-free complete allocation may not always exist. Our first objective is to find an envy-free allocation of maximum size (see Example 1). In other words, among the set of all envy-free allocations we find an allocation that maximizes the number of assigned agents (or houses when $m \leq n$).

In their paper [9] describe an efficient polynomial time algorithm to find an envy-free allocation in ordinal instances. However, they restrict themselves to complete envy-free allocations, i.e., where each agent is assigned a house. Their algorithm returns 'empty' when the number of available houses falls below the number of agents. We relax this constraint and return a maximum size envy-free allocation which need not be complete on both binary and weighted instances.

The key difference is that we allow allocations along zero edges (aka houses that are valued zero and are not adjacent). We start by discussing binary instances.

Proposition 1 ([2]). *Every bipartite graph $G = (N \cup H, E)$ admits a unique partition of $N = N_S \cup N_L$ and of $H = H_S \cup H_L$ such that every envy-free matching in G is contained in $G[N_L, H_L]$, a maximum matching in $G[N_L, H_L]$ is a maximum envy-free matching in G, and it can be computed in polynomial time.*

Proposition 2. *Given a binary instance, an envy-free allocation of maximum size can be computed in polynomial time.*

Proof sketch. Given a binary instance $\langle N, H, V \rangle$, we start by constructing a bipartite graph $G = (N \cup H, E)$ such that given $i \in N$ and $h \in H$, $(i, h) \in E$ if and only if $v_i(h) = 1$. We do the following:

(1) Find a envy-free matching M by using Proposition 1. Add all (agent, house) pairs in M to A.
(2) While there exists an unassigned (agent, house) pair (j, h) in A such that each agent i with valuation $v_i(h) > 0$ is matched in M, we add the pair (j, h) to A.

By Proposition 1, M is envy-free. Step (2) of our algorithm allocates houses in H_L as long as there is an unassigned agent. Thus, A is a maximum size EF allocation. □

The above algorithm is based on binary bipartite matchings, and thus, it fails to work when we allow for more expressive preferences. We develop a polynomial time algorithm to find maximum size allocations with zero envy on weighted instances below (Algorithm 1).

Algorithm description. Algorithm 1 creates a bipartite graph where each agent is only adjacent to houses that are its most preferred. In other words, if there is an edge between an agent i and a house h, then $v_i(h)$ is positive and maximum among all houses remaining in the graph. The algorithm proceeds by then removing all houses in any inclusion-minimal Hall violators. This process repeats by updating the highest valued house among the remaining houses for each agent, and adding, or retaining, the corresponding edges. We denote the complement graph of G as $\overline{G}$ and $\overline{G} - M$ denotes the graph we get by deleting all edges and vertices of M from $\overline{G}$.

Algorithm 1 A maximum size envy-free allocation

Input: A house allocation instance $\langle N, H, V \rangle$
Output: A maximum size envy-free allocation
1: **for** each agent $i \in N$ **do**
2: Let $h_i^{\max} \in \text{argmax}_{h \in H}\, v_i(h)$
3: **end for**
4: Let $E_{\max} := \{(i,h) | i \in N, h \in H, v_i(h) = v_i(h_i^{\max}), v_i(h) > 0\}$
5: Create a bipartite graph $G = (N \cup H, E_{\max})$
6: **if** there exists an inclusion-minimal Hall violator (N', H') in G **then**
7: Delete H' from H; *removing houses that cause envy.*
8: **else**
9: **return** Allocation A = union of maximum size matching M in G and a maximum size matching in $\overline{G} - M$.
10: **end if**
11: Go to line 1

After either all houses are considered or no Hall violator is found, the algorithm returns a *maximum size allocation* which is produced by union of a maximum size matching M in the induced graph G and a maximum size matching in the graph $\overline{G} - M$. The latter allows us to assign houses to agents that do not want them, i.e., $v_i(h) = 0$, without creating envy.

The next lemma gives a natural invariant for the algorithm.

Lemma 1. *Given a weighted instance, any house removed by Algorithm 1 cannot be a part of any envy-free allocation.*

Theorem 1. *Given a weighted instance, Algorithm 1 returns an envy-free maximum size allocation in polynomial time.*

Proof. First, note that Algorithm 1 runs in polynomial time because every component of the algorithm including finding a inclusion-minimal Hall violator [9, 2] and computing a maximum size bipartite matching runs in time polynomial in n and m. Therefore, it suffices to prove that (i) every house removed by the algorithm cannot be contained in any envy-free allocation, (ii) no assigned house causes an envy, and (iii) a maximum size bipartite matching returns a maximum size allocation among all envy-free allocations.

Statement (i) immediately follows from Lemma 1 and the fact that each agent's valuations for the houses unassigned in M is zero. Statement (ii) follows from the fact we observe that every agent only has edges to its most preferred house among the houses retained, no further edges can be added, and no Hall violators exist. Moreover, the houses that are assigned in $\overline{G} - M$ do not cause envy, since no assigned agent has a higher valuation for them, and no unassigned agent has a positive valuation for them. Statement (iii) follows from the observation that Algorithm 1 finds a maximum size matching in the induced graph G. The unassigned houses and unassigned agents in M are assigned in a maximum matching in $\overline{G} - M$. Thus, no agent or house that is already assigned in M is reassigned. Moreover, since we find a maximum matching in $\overline{G} - M$, it assigns maximum number of agents to the houses they value zero. Thus, the algorithm returns the maximum size envy-free allocation on the instance. □

4 Utilitarian Welfare

In this section, we show that fair allocations can be obtained efficiently by introducing a utilitarian welfare constraint and show the connection between complete allocations and welfare maximizing ones when restricting the number of houses.

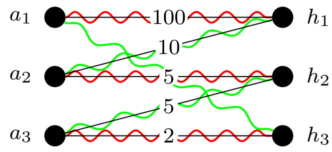

Figure 2: The min #envy max USW allocation (shown in red) is not a min #envy complete allocation (shown in green) in weighted instances.

We begin the section by designing an EF allocation with maximum welfare.

Proposition 3. *Given a weighted instance, an envy-free maximum* USW *allocation, should it exist, can be computed in polynomial time.*

Note that while a maximum size envy-free allocation can be computed in polynomial time, finding a complete allocation that minimizes the number of envious agents is NP-hard [16] (see example given in Figure 1).

4.1 Minimum #Envy

We start by showing that finding a USW-maximizing allocation that minimizes #envy can be done in polynomial time. The algorithm constructs a bipartite graph and uses minimum cost perfect matching [21]. The details of the construction and proofs are relegated to [12].

Theorem 2. *Given a weighted instance, a min #envy max* USW *allocation can be computed in polynomial time.*

While finding a min #envy complete allocation is shown to be NP-hard even for binary instances [16], we establish a relation between min #envy complete and min #envy max USW allocations when $m \leq n$ by extending a minimum #envy maximum USW allocation to a complete allocation.[4]

Proposition 4. *In a binary instance when $m \leq n$, given a min #envy max* USW *allocation A, a complete allocation $\hat{A}$ can be constructed in polynomial time such that*

(i) A and $\hat{A}$ have equal USW *and #envy, and*
(ii) $\hat{A}$ has minimum #envy among all complete allocations.

Note that Proposition 4 does not hold for weighted instances (4).

Example 2 (Proposition 4 does not hold for weighted instances)**.** *Consider three houses and three agents as shown in Figure 2. The values are shown on edges. The min #envy max* USW *allocation shown in red must allocate h_1 to a_1 with the value of* 100 *to satisfy the maximum* USW *constraint. This allocation results in creating two envious agents (agents 2 and 3). However, The min #envy complete allocation (shown in green) has exactly one envious agent.*

Nevertheless, we show that a min #envy complete allocation can be computed in polynomial time for weighted instances when $m \leq n$. The proof is relegated to the full version of this paper [12].

Remark 1. *First, given a maximum size envy-free allocation we cannot 'append' it to achieve a min #envy complete allocation even for binary instances. In Example 1, the allocation indicated by orange lines cannot be simply completed to reach the min #envy complete allocation (shown in green). Second, when $m > n$ a min #envy max*

[4] For binary instances, Madathil et al. [19] independently showed that a min #envy complete allocation (termed as "optimal" house allocation) can be computed in polynomial time when $m \leq n$.

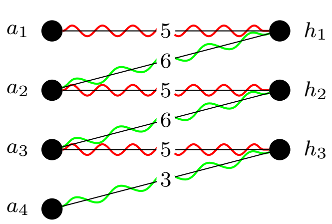

Figure 3: An example depicting the difference between min #envy and min total envy in USW-maximizing allocations. Red has three envious agents, namely, a_2, a_3, a_4 and green has one envious agent a_1 but both have total envy five.

USW *allocation may leave more agents envious compared to a min #envy complete allocation. In Example 1, the min #envy max* USW *allocation indicated by blue leaves two agents envious while the min #envy complete allocation (shown in green) only leaves one agent envious.*

Theorem 3. *Given a weighted instance, when $m \le n$, a min #envy complete allocation can be computed in polynomial time.*

Our approach to find a min #envy max USW allocation heavily relies on the set of USW maximizing allocations. Next, we discuss two observations about the necessity of utilizing USW maximizing allocations and restriction on the number of houses.

4.2 Minimum Total Envy

When focusing on total envy of agents, under binary valuations, the total envy can be seen as the number of distinct envy relations between all pair of agents. Whereas in weighted instances, individual values for each assigned houses (and not just pairwise relations) play an important role in computing the total envy.

Example 3. *Consider the weighted instance given in Figure 3. There are two allocations that maximize the* USW *with a total welfare of* 15. *One allocation (shown in red) leaves three agents envious (namely, a_2, a_3, and a_4) with a total envy of* 5*; while another allocation (shown in green) only contains one envious agent (a_1) and still generates the total envy of* 5.

This example illustrates that the proposed algorithms for finding min #envy max USW cannot be readily used for finding min total envy max USW. In [12], we present an example (Example 4) which shows that this challenge persists even for binary instances. Nonetheless, we show that one can achieve a min total envy max USW allocation in polynomial time by constructing a bipartite graph, similar to the algorithm for min #envy max USW, with a carefully crafted cost function that encode envy and USW as cost, and utilizing algorithms for finding a minimum cost perfect matching.

Algorithm description. The algorithm (Algorithm 6 in[12]) proceeds by constructing a bipartite graph $G = (N \cup H', E)$ where the set H' is constructed by adding a set of n dummy houses to the set of houses H. That is, $H' = H \cup \{h^i \mid i \in N\}$. Given an agent $i \in N$ and $h \in H'$, we have $(i, h) \in E$ if and only if $h \in H$ and $v_i(h) > 0$, or $h \in H' \setminus H$. We construct a cost function c on edges of G. Before defining the cost, for ease of analysis, we multiply each of the values in the valuation V with an appropriate scalar such that the following holds: for each agent i and house h, if $v_i(h) > 0$, then $v_i(h) \ge 1$. Now we define two components of the cost function c, namely, envy component c_{envy} and USW component c_{sw} such that $c = c_{envy} + c_{sw}$. Let $H_i^{\max}$ be the set of most preferred houses in H for agent i, i.e., $H_i^{\max} = \text{argmax}_{h \in H}\, v_i(h)$. For ease of exposition we assume $v_i(h) = 0$ for each agent i and a dummy house $h \in H' \setminus H$. For an edge $(i, h) \in E$, if $h \in H_i^{\max}$, then define $c_{envy}(i, h) = 0$; otherwise $c_{envy}(i, h) = \sum_{\bar{h} \in H} \max\{v_i(\bar{h}) - v_i(h), 0\}$. We denote $\sum_{i \in N} \sum_{\bar{h} \in H} v_i(\bar{h}) + 1$ by L. Furthermore, we define the USW component of the cost c for an edge $(i, h) \in E$ as $c_{sw}(i, h) = -v_i(h) \cdot L$. Finally, we return a minimum cost perfect matching matching in G.

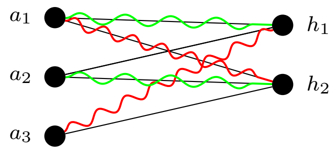

Figure 4: Two complete allocations with different total envy. Total envy of the red allocation is two due to agent a_2, the same for green is one due to a_3.

Theorem 4. *Given a weighted instance, a min total envy max* USW *allocation can be computed in polynomial time.*

Proof Sketch. Suppose the algorithm (Algorithm 6 in [12]) returns allocation A. Then A corresponds to a minimum cost perfect matching in the constructed graph G. We show that by minimizing the cost, we maximize the welfare and minimize the total envy. Observe that the cost of each pair $(i, h) \in A$ has two components, namely, USW component $-v_i(h) \cdot L$ and a envy component. To complete the proof we show that (i) since L is large, a minimum cost matching in G maximizes USW; (ii) the envy component of the cost of a perfect matching correctly computes the total envy of each agent. It follows from the fact that in a max USW allocation every house valued higher than house h by agent i must be allocated for each $(i, h) \in A$. ☐

We show a result analogous to Proposition 4 holds for min total envy even for weighted instances.

Proposition 5. *Given a weighted instance with $m \le n$, let A^* be a min total envy max* USW *allocation. Then a complete allocation A can be constructed in polynomial time such that*

(i) A^* and A has equal USW *and envy, and*
(ii) A is a min total envy complete allocation.

Corollary 1 follows immediately from Proposition 5.

Corollary 1. *Given a weighted instance, a min total envy complete allocation can be computed in polynomial time when $m \le n$.*

Observe that even when $m \le n$, there may be several complete matchings with different total envy. In Figure 4, both matchings are complete because they assign all the houses, however, the allocation indicated by red yields a higher total envy (two by a_2) than the green one (one by a_3).

5 Egalitarian Welfare

When the efficiency measure is maximizing the utilitarian welfare, any maximum-weight matching on the induced bipartite graph can find such an allocation in polynomial time. However, the problem of finding an allocation that maximizes the egalitarian welfare has received less attention in the house allocation setting.

While in fair division finding an allocation that maximizes the egalitarian welfare is NP-hard[5], we show that in the house allocation setting wherein agents are restricted to receive at most one house, an egalitarian solution can be found in polynomial time.

[5] When agents can receive multiple items, an egalitarian allocation always exists but computing a egalitarian allocation is NP-hard [6].

Algorithm 2 Finding an allocation of max ESW

Input: A house allocation instance $\langle N, H, V\rangle$.
Output: An allocation with maximum egalitarian welfare.
for $k = n$ to 1 **do**
  Let $v_{\max} := \max\{v_i(h) > 0 \mid i \in N, h \in H\}$ and $v_{\min} := \min\{v_i(h) > 0 \mid i \in N, h \in H\}$
  ▷ LOOPING ALL UNIQUE VALUES OF HOUSES IN H
  for $\beta = v_{\max}$ to $v_{\min}$ **do**
    Create a bipartite graph $G_\beta = (N \cup H, E)$ s.t. there exists an edge between each agent $i \in N$ and each house $h \in H$ if $v_i(h) \geq \beta$
    Let $A :=$ a maximum size matching of G_β
    if there exist k allocated agents in A, **return** Allocation A
  end for
end for

Note that in binary instances, finding an egalitarian allocation is equivalent to finding an envy-free allocation of maximum size (Proposition 2) since in every allocation the egalitarian welfare is either zero or one. When it comes to weighted instances, however, the goal is to maximize the number of agents who receive a positive value and, conditioned on that, maximize the value of the worst-off agent. We use this intuition to search for an allocation that maximizes the number of agents that receive a positively valued house.

Algorithm description. The algorithm (Algorithm 2) begins by considering the maximum number of agents $k = n$ who can potentially receive a positively-valued house. Consider the set of all agent valuations in decreasing order. For each such positive value $\beta > 0$, create a bipartite graph G_β where there is an edge between any agent-house (i, h) pair if $v_i(h) \geq \beta$. Now, find a maximum-size matching M on G_β. If the size of the matching is k, the algorithm returns M as the required allocation. Otherwise, it repeats this process by decreasing the size to $k = k - 1$.

Theorem 5. *Given a weighted instance, an egalitarian welfare maximizing allocation can be found in polynomial time.*

Proof. Suppose that Algorithm 2 returns an allocation A of size k where every allocated agent receives positive value of at least β. To prove the theorem we need to show that (i) k is the largest number of agents that can simultaneously receive positive value and (ii) $\mathsf{ESW}(A)$ is maximum among all allocations of size at least k. Note that it is sufficient to prove this for size exactly k since we cannot increase ESW by increasing the size of A. Since A is a maximum size matching in G_β, from definition of G_β, it holds that we cannot allocate more than k agents to the houses they value at least β. Furthermore, the size k decreases from n, and the algorithm returns the first k-sized allocation where each assigned agent receives positive value. Thus A is the largest possible allocation where each assigned agent gets some positive value since we iterate over all positive values of β. Thus we show (i). Moreover, since we start by setting β to the highest possible value of ESW and decrease step by step, there does not exist a k size allocation for a higher value of β. Therefore, for any k-sized allocation β is the maximum ESW. Thus we show (ii). Note that the possible values of β are bounded by the distinct values agents have towards the houses. There are $O(mn)$ values that can be assumed by β. Thus the algorithm runs in polynomial time. □

Clearly, a maximum egalitarian allocation may not be unique. Thus, a natural question is whether we can find a fair allocation among all such allocations. We first show that analogous to its utilitarian counterpart (Proposition 3), an envy-free allocation (if one exists) of maximum ESW can be computed in polynomial time.

Algorithm 3 Finding an EF allocation of maximum ESW

Input: A house allocation instance $\langle N, H, V\rangle$.
Output: An EF allocation with maximum egalitarian welfare.
Let k and β denote the size and ESW of an allocation returned by Algorithm 2.
Create valuation V' s.t $v'_i(h) = v_i(h)$ if $v_i(h) \geq \beta$; $v'_i(h) = 0$ otherwise, for an agent $i \in N$ and house $h \in H$.
Let A be the allocation returned by Algorithm 1 given $\langle N, H, V'\rangle$.
if there exist k allocated agents in A, **return** Allocation A
return $\emptyset$

Algorithm description. Given an instance $I = \langle N, H, V\rangle$, we find the max ESW allocation with welfare at least β for k agents using Algorithm 2. Then we find an envy-free allocation, if exists, with egalitarian welfare at least β for k agents. We construct a reduced valuation V' where $v'_i(h)$ is set to $v_i(h)$ if $v_i(h) \geq \beta$ and is zero otherwise for an agent $i \in N$ and house $h \in H$. This is to ensure that we satisfy ESW. Then, we invoke Algorithm 1 as a subroutine to find a EF allocation A in $\langle N, H, V'\rangle$. If k agents receive value at least β in A, then we return the allocation A; otherwise we return an empty allocation since no EF allocation of max ESW exists.

Theorem 6. *Given a weighted instance, an envy-free allocation of maximum* ESW*, should it exist, can be computed in polynomial time.*

Proof. We have the following property of the allocation A returned by Algorithm 3. Each agent that is assigned a house in A receives value at least β. The statement follows from the facts that each agent is assigned a house that it values positively by Algorithm 1 and each positive value in V' is at least β. However, there may exist agents that are not assigned any house in A. Thus, if Algorithm 1 returns an allocation that does not assign k agents, from the definition of maximum egalitarian welfare we conclude that there is no EF allocation of max ESW. The correctness of this step follows from Theorem 1. It shows that in A, maximum number of agents are assigned with positive value. So an agent i that is unassigned in A cannot be assigned to a house it likes without generating envy. Therefore, there is no EF allocation that can match i to a house that it values β or more. Using Theorem 1, we have that allocation A is EF. This completes the proof of correctness. Since Algorithms 1 and 2 run in time polynomial in n and m, Algorithm 3 runs in polynomial time. □

5.1 Minimum #Envy

We aim to find an ESW welfare maximizing allocation that minimizes the number of envious agents. Under binary valuations, the ESW is either zero or one. When the ESW is one, we return a complete, envy-free allocation. For ESW zero, an empty allocation is the optimal solution. In contrast to the utilitarian welfare setting (Theorem 2), finding a max ESW allocation that is min #envy is intractable in a weighted instance.

Theorem 7. *Given a weighted instance, finding a min #envy max* ESW *allocation is* NP-hard.

proof sketch. We prove this by showing a reduction from the problem of finding a min #envy complete allocation that is known to be NP-complete [16]. Given an instance $I = \langle N, H, V\rangle$ of the minimum #envy complete problem, we build an equivalent instance of the min #envy max ESW problem. For each agent $i \in N$ and house $h \in H$, we create the valuation $v'_i(h)$ by adding a positive small

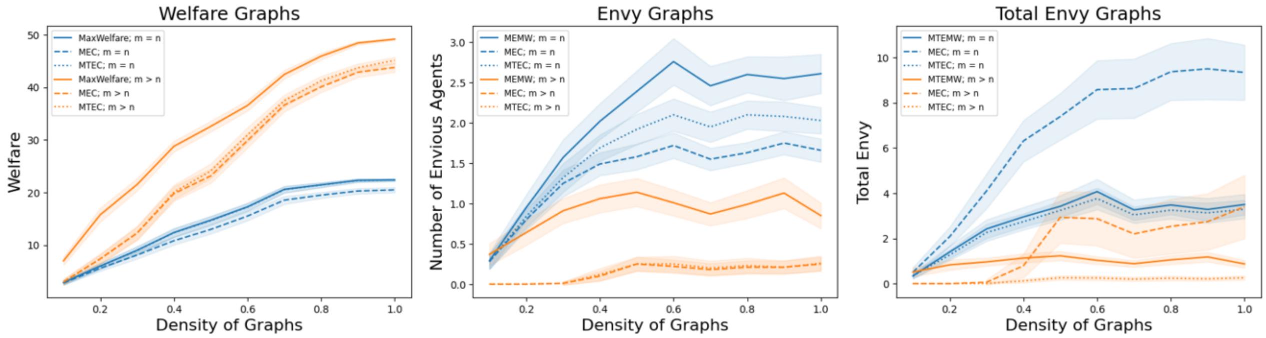

Figure 5: The averages of #envy (number of envious agents), total envy, and USW over 100 random trials for each graph density. Codes in the legend refer to MTEC: min total envy complete; MEC: min #envy complete; MEMW: min #envy max USW; MTEMW: min total envy max USW.

value β to the valuation $v_i(h)$. Thus, all max ESW allocations assign a positively valued house to each agent under the new valuation. We show the hardness of the problems lies in minimizing #envy under the completeness requirement. In [12] we show the equivalence of the two instances to complete the proof. □

It is easy to check that the decision version of min #envy max ESW - where we check if there exists a max ESW allocation with #envy at most t - is in NP. Thus the problem is NP-complete.

5.2 *Minimum Total Envy*

While minimizing the total amount of envy experienced by the agents, restricting the search space to the maximum egalitarian welfare allocations does not result in any computational improvement. The problem remains computationally as hard as finding a min total envy complete allocation.

Theorem 8. *Given a weighted instance, finding a min total envy of max* ESW *is as hard as finding a min total envy complete allocation.*

The proof uses the same construction as in Theorem 7. We defer the details to the full version of this paper [12]. If in return the objective is to minimize the maximum total of envy experienced by agents (minimax total envy) the problem becomes NP-hard. Note that similar to Theorem 7, in the decision version of the problem, one can check whether minimax total envy of the allocation is at most t, implying that the problem is NP-complete.

Theorem 9. *Given a weighted instance, finding a minimax total envy max* ESW *allocation is* NP*-hard.*

The detailed proof can be found in [12]. In a nutshell, we provide a reduction from the Independent Set problem in cubic graph [10]. Note that even though our construction is similar to [19]'s hardness reduction for finding a minmax total envy complete allocation[6], our reduction further ensures that every agent receives a positive value.

6 Experiments

We experimentally investigate the welfare loss and fairness of the proposed algorithms on randomly generated bipartite graphs. For a fixed number of agents, we varied the number of houses ($m = n$ and $m > n$) and considered both binary and weighted valuation functions V. We modelled preferences by iterating over the density of edges ($\lambda \in [0.1, 1.0]$) in the corresponding bipartite graph. For each instance, defined by (m, λ, V), we ran 100 trials, on a randomly generated graph G that satisfied the (m, λ, V) constraints. We compared the maximum USW achieved by min #envy max USW or min total envy max USW with the welfare of min #envy complete, and min total envy complete allocations to understand the price of fairness. Next, we compared the #envy (and total envy) of min #envy complete (resp. min total envy complete) with that of the min #envy max USW and min total envy complete (resp. min total envy max USW and min #envy complete). The plots above show us the average value of these metrics over the 100 trials for weighted valuations, and highlight the 95%-Confidence Interval.[7]

Observations. When there is an abundance of houses, the envy and total envy of all allocations decreases and the USW increases. Similarly, as the graph grows denser (i.e. $\lambda > 0.4$), welfare increases and, under binary valuations, envy and total envy vanish. For weighted valuations too, we notice a slight decrease, but they still persist. As expected, the number of envious agents in a min #envy complete allocation is least, followed by min total envy complete and min #envy max USW. The lower USW of min total envy complete can be attributed to leaving highly valued houses unallocated. Notably, min #envy complete has higher total envy than min total envy max USW, since it would prefer one highly envious agent to multiple slightly envious ones. Additional plots and discussions on experiments can be found in [12].

7 Concluding Remarks

Our investigation on different efficiency and fairness concepts gives rise to several intriguing open questions. For example, the computational complexity of minimizing total envy remains unsolved. Moreover, one can ask if we can guarantee approximations of welfare to achieve EF or relaxations of EF; or whether the complexity of the problems change when considering strict ordinal preferences, Borda valuations, or pairwise preferences.

Acknowledgments

Hadi Hosseini acknowledges support from NSF IIS grants #2144413 and #2107173. We thank the anonymous reviewers for their helpful comments.

[6] Madathil et al. [19] refer to this problem as "egalitarian house allocation".

[7] The source code is publicly available at https://github.com/medha-kumar/DegreeOfFairnessInEfficientHouseAllocation.git.

References

[1] A. Abdulkadiroğlu and T. Sönmez. School choice: A mechanism design approach. *American economic review*, 93(3):729–747, 2003.

[2] E. Aigner-Horev and E. Segal-Halevi. Envy-free matchings in bipartite graphs and their applications to fair division. *Information Sciences*, 587: 164–187, 2022.

[3] N. Bansal and M. Sviridenko. The santa claus problem. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of Computing*, pages 31–40, 2006.

[4] K. Belahcène, V. Mousseau, and A. Wilczynski. Combining fairness and optimality when selecting and allocating projects. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI-21)*, pages 38–44. International Joint Conferences on Artificial Intelligence Organization, August 2021. doi: 10.24963/ijcai.2021/6. URL https://doi.org/10.24963/ijcai.2021/6.

[5] N. Benabbou, M. Chakraborty, E. Elkind, and Y. Zick. Fairness towards groups of agents in the allocation of indivisible items. In *The 28th International Joint Conference on Artificial Intelligence (IJCAI'19)*, 2019.

[6] S. Bouveret and M. Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 30(2):259–290, 2016.

[7] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum nash welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–32, 2019.

[8] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaitre, N. Maudet, J. Padget, S. Phelps, J. A. Rodriguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30(1):3–32, 2006.

[9] J. Gan, W. Suksompong, and A. A. Voudouris. Envy-freeness in house allocation problems. *Mathematical Social Sciences*, 101:104–106, 2019.

[10] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

[11] H. Hosseini, J. Payan, R. Sengupta, R. Vaish, and V. Viswanathan. Graphical house allocation. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, pages 161–169, 2023.

[12] H. Hosseini, M. Kumar, and S. Roy. The degree of fairness in efficient house allocation. *arXiv preprint arXiv:2407.04664*, 2024.

[13] H. Hosseini, A. McGregor, R. Sengupta, R. Vaish, and V. Viswanathan. Tight approximations for graphical house allocation. In *Proceedings of the 2024 International Conference on Autonomous Agents and Multiagent Systems*, page forthcoming, 2024.

[14] A. Hylland and R. Zeckhauser. The efficient allocation of individuals to positions. *Journal of Political economy*, 87(2):293–314, 1979.

[15] N. Kamiyama. The envy-free matching problem with pairwise preferences. *Information Processing Letters*, 172:106158, 2021.

[16] N. Kamiyama, P. Manurangsi, and W. Suksompong. On the complexity of fair house allocation. *Operations Research Letters*, 49(4):572–577, 2021.

[17] J. K. Lenstra, D. B. Shmoys, and É. Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Mathematical Programming*, 46:259–271, 1990.

[18] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 125–131, 2004.

[19] J. Madathil, N. Misra, and A. Sethia. The complexity of minimizing envy in house allocation. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems, AAMAS '23, London, UK, May 29-June 2, 2023*, 2023.

[20] T. Mashiat, X. Gitiaux, H. Rangwala, P. Fowler, and S. Das. Trade-offs between group fairness metrics in societal resource allocation. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 1095–1105, 2022.

[21] L. Ramshaw and R. E. Tarjan. On minimum-cost assignments in unbalanced bipartite graphs. *HP Labs, Palo Alto, CA, USA, Tech. Rep. HPL-2012-40R1*, 20, 2012.

[22] A. E. Roth. Incentive compatibility in a market with indivisible goods. *Economics letters*, 9(2):127–132, 1982.

[23] A. E. Roth, T. Sönmez, and M. U. Ünver. Kidney exchange. *The Quarterly journal of economics*, 119(2):457–488, 2004.

[24] L. Shapley and H. Scarf. On cores and indivisibility. *Journal of mathematical economics*, 1(1):23–37, 1974.

[25] D. B. Shmoys and É. Tardos. An approximation algorithm for the generalized assignment problem. *Mathematical Programming*, 62(1-3):461–474, 1993.

[26] L.-G. Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16:557–567, 1999.