Local Envy-Freeness in House Allocation Problems

Aurélie Beynier, Yann Chevaleyre, Laurent Gourvès, Ararat Harutyunyan, Julien Lesca, Nicolas Maudet, Anaëlle Wilczynski

Noname manuscript No.
(will be inserted by the editor)

Local Envy-Freeness in House Allocation Problems

**Aurélie Beynier · Yann Chevaleyre ·
Laurent Gourvès · Ararat Harutyunyan ·
Julien Lesca · Nicolas Maudet ·
Anaëlle Wilczynski**

Received: date / Accepted: date

Abstract We study the fair division problem consisting in allocating one item per agent so as to avoid (or minimize) envy, in a setting where only agents connected in a given network may experience envy. In a variant of the problem, agents themselves can be located on the network by the central authority. These problems turn out to be difficult even on very simple graph structures, but we identify several tractable cases. We further provide practical algorithms and experimental insights.

Keywords Object allocation · Envy-Freeness · Complexity · Algorithms.

1 Introduction

Fairly allocating resources to agents is a fundamental problem in economics and computer science, and has been the subject of intense investigations [17, 23]. Recently, several papers have explored the consequences of assuming in such settings an underlying network connecting agents [2, 6, 11, 19, 23]. The most intuitive interpretation is that agents have limited information regarding the overall allocation. Two agents can perceive each other if they are directly connected in the graph.

A. Beynier, N. Maudet
Sorbonne Université, CNRS,
Laboratoire d'Informatique de Paris 6, LIP6,
F-75005 Paris, France
E-mail: {aurelie.beynier,nicolas.maudet}@lip6.fr

Y. Chevaleyre, L. Gourvès, A. Harutyunyan, J. Lesca, A. Wilczynski
Université Paris-Dauphine, PSL, CNRS,
LAMSADE, Paris, France
E-mail: {yann.chevaleyre,laurent.gourves,ararat.harutyunyan
julien.lesca,anaelle.wilczynski}@dauphine.fr

A fairness measure, very sensitive to the information available to agents, is the notion of *envy* [31]. Indeed, envy occurs when an agent prefers the share of some other agents over her own. Accounting for a network topology boils down to replacing "other agents" by "neighbors". The notion of envy can thus naturally be extended to account for the limited visibility of the agents. Intuitively, an allocation will be *locally envy-free* if none of the agents envies her neighbors. This notion has been referred as *graph*, *social*, or *local envy-freeness* [2, 6, 19, 22, 23, 30]. It finds its origins in Festinger's work on social comparisons which are not made globally but locally, i.e. with respect to an individual's neighbors in the social network [29].

In this paper, we are concerned with the allocation of indivisible goods within a group of agents. The setting we study in this paper is arguably one of the simplest in resource allocation, known in economics as *house allocation* [1, 36, 49]: agents have (strict) preferences over items, and each agent must receive exactly one item. In the case of a complete network, envy-freeness is not a very exciting notion in that setting. Indeed, for an allocation to be envy-free, each agent must get her top object (and this is obviously also a Pareto-optimal allocation in that case). When an agent is only connected to a subset of the other agents, she may not need to get her top-resource to be envy-free. The locations of the resources on the graph as well as the connections between the agents are then crucial issues in order to compute a locally envy-free allocation.

To see how the network can make a difference, consider the following scenario.

Example 1 Suppose for instance a team of workers taking their shifts in sequence, to which a central authority must assign different jobs. Workers have preferences regarding these jobs. As the shifts are contiguous and as the employees work at the same place, they have the opportunity to see the job allocated to some other workers, as one ends and the other one begins her shift. This would be modeled as a line topology in our setting as depicted on the graph in Figure 1. To make things concrete, suppose there are three jobs, chop the tree, mow the lawn, and trim the hedge, and three gardeners (1, 2 and 3) with preferences $1 : chop \succ mow \succ trim$, $2 : mow \succ chop \succ trim$, $3 : chop \succ trim \succ mow$, taking shifts in order 1, 2 and finally 3. On Figure 1, rankings are mentioned over agents (with top jobs at the top, etc.)

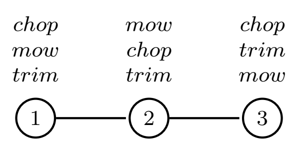

Fig. 1: Example of working locations and preferences of three gardeners over three jobs to perform

By allocating the job *chop the tree* to agent 1, *mow the lawn* to agent 2, and *trim the hedge* to agent 3, we get an envy-free allocation if we disregard the fact that agent 3 may be envious of agent 1. Note that a locally envy-free allocation is not necessarily Pareto-optimal (take the same allocation, but the ranking of agent 1 to be $trim \succ chop \succ mow$). However, giving her top item to each agent if possible will always be an envy-free Pareto-optimal allocation in *any* network.

Now, consider that we switch the locations of agent 2 and agent 3, as depicted in the graph of Figure 2.

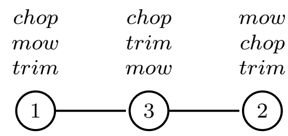

Fig. 2: Switch in the working locations of the three gardeners

In this instance, there is no locally envy-free allocation. In fact, if agent 3 gets the job *chop the tree*, she will be envied by agent 1. If agent 1 or agent 2 gets the job *chop the tree*, agent 3 will be envious of one of her neighbors. From these observations, it is worth investigating the problem of deciding whether a locally envy-free allocation exists.

When a locally envy-free allocation does not exist, one can try to optimize the number of agents that are locally envy-free or to optimize the degree of non-envy in the network. Hence, by allocating the job *mow the lawn* to agent 1, *trim the hedge* to agent 3 and *chop the tree* to agent 2, only one agent is envious of another agent (agent 3 is envious of agent 2). This solution both optimizes the number of agents that are locally envy-free and the degree of non-envy in the network. In this paper, we will also investigate these optimization problems.

The reader may object that, in the first network of Example 1, agent 3 may still be envious of agent 1, because she knows that this agent must have received the task agent 2 didn't get, i.e. *chop the tree*. This is a valid point, to which we provide two counter-arguments. First, as a technical response, note that in general agents would not know exactly who gets the items they do not see. Thus, although agents may know that they must be envious of *some* agents, they cannot identify which one, which makes a significant difference in the case of envy. It could also be that agents actually do not know which objects are to be allocated in the first place. For instance, while the central authority may know the preferences of gardeners over all the possible tasks to be performed, the gardeners themselves may not know each morning exactly which task is to be performed on that day. Our second point is more fundamental and concerns the model and the motivation of this work. Clearly, the existence of a network may be due to an underlying notion of proximity (either geographical, or temporal as in our example) in the problem. However, another interpretation of the meaning of links must be emphasized: links may represent envy the central authority is concerned with. In other words, although there may theoretically be envy among all agents, the central authority may have reasons to only focus on some of these envy links. For instance, you may wish to avoid envy among members of the same team in your organization, because they actually work together on a daily basis (in that case links may capture team relationships). Under this interpretation, a network of degree $n-2$ (the total number of agents minus 2), where an agent can envy everyone except one another agent, could for instance model a situation where agents team-up in pairs and conduct a task together, sharing their resources. In a similar vein, we may focus on avoiding envy among "similar" agents, because they may be legitimate to complain if they are not treated equally despite similar competences, for instance.

1.1 Related work

Our work is connected to a number of recent contributions addressing fair allocation on graphs.

Both Abebe et al. [2] and Bei et al. [11] studied envy-freeness and proportionality for the *cake cutting problem* where comparisons between agents are limited by an underlying network structure. Cake cutting deals with the fair allocation of divisible goods (e.g. land) while the present work is devoted to indivisible resources.

Bredereck et al. [19] introduced a model with indivisible resources which is very close to ours. The underlying graph is directed (agent u can envy agent v if the arc oriented from u to v exists), and the number of objects that an agent receives is not fixed (it can be 0) and it may differ between the agents. The present work deals with undirected graphs, and every agent must receive exactly one object. An instance described with an undirected graph can also be described with a directed graph because an edge can be replaced by two arcs of opposite orientations. However, an algorithm or a reduction designed for a kind of graphs (directed or not) may not translate to the other kind. Bredereck et al. investigate the standard and parameterized computational complexity of finding an allocation where the agents have additive and monotone utility functions over the objects. Envy-freeness has to be satisfied along the arcs of the directed graph, together with an additional requirement which can be *completeness* (all the objects are assigned), *Pareto-efficiency*, or the fact that the *utilitarian social welfare* is maximized.

There exist different fairness criteria (max-min fair share, proportionality, envy-freeness, CEEI, etc.), which are connected by implication relations and form a *scale of fairness* according to the strength of their requirement [18]. For example, under mild assumptions on the agents' utilities, envy-freeness implies proportionality which implies max-min fair share. The relations between these fairness concepts were recently enriched by Aziz et al. with a novel notion called *graph epistemic envy-freeness* [6]. Agents are solely aware of the shares of their neighbors in a given social network (a directed graph). An agent i is envious if her share $A(i)$ is worse than a neighbor's share, or any allocation of the objects not present in the shares that she is aware of, must contain a share that agent i finds better than $A(i)$.

Recently, house allocation settings have been discussed, notably in relation with swap dynamics [24, 34]. In particular, Gourvès et al. [34] show how graph structures can affect the complexity of some decision problems regarding such dynamics. More specifically, they study the complexity of deciding whether some object is reachable by a given agent or whether some allocation is reachable by a sequence of swaps among agents. The complexity of searching a Pareto-efficient solution is also studied. The reachable object problem is re-examined by Saffidine and Wilczynski [44], assuming that the number of swaps and the total duration of the process are limited. Even more recently, Kondratev and Nesterov unveiled surprising connections in house allocation settings between the minimization of the number of envious agents, and popular matchings [40]. Their notion of envy slightly differs from ours though, in the sense that it excludes envy towards those agents who get their preferred item (which they call "inevitable" envy).

The allocation *of* a graph has also recently been studied [13, 16, 37, 46]. In this context, the nodes of the graph represent indivisible resources to allocate and edges formalize connectivity constraints between the resources: each agent must receive items which form a connected component in the graph. The graph structure enables to capture dependencies between the resources, like spatial dependencies for pieces of land or time constraints.

In a similar framework, some computational aspects of allocating agents on a line are discussed by Aziz et al. [7]. In this setting, the line concerns slots to allocate to the agents, and can be viewed as the problem of placing the agents at the nodes of a line. The agents have specific target locations on the line, which induces a domain restriction (stronger than single-peakedness).

Several ways for a central authority to control fair division have been discussed by Aziz et al. [8]: the structure of the allocation problem can be changed by adding or removing items to improve fairness. Interestingly our model introduces a new type of control action: locating agents on a graph. Finally, because envy-freeness cannot be guaranteed in general (with indivisible items), and as related decision problems can be difficult even in simple settings [38, 41], different notions of degree (or relaxation) of envy have been studied [13, 21, 27, 41, 42], and the relation between some of these relaxations has been studied by Amanatidis et al. [3].

1.2 Contributions and organization

A formal definition of the model, together with the definition of the main problems that we address, are provided in Section 2. Section 3 is dedicated to the problem, called DEC-LEF, of deciding if a central planner, who has a complete knowledge of the social network and the agents' rankings of the objects, can allocate the objects such that no agent will envy a neighbor. Note that in this setting the central planner does not decide where the agents are located on the network. We identify intractable and tractable cases of this decision problem, with respect to the number of neighbors of each agent, that is the degree of the nodes in the graph representing the social network. Remarkably, we show that the problem turns out to be intractable even on social networks with simple structure: when agents are matched one-to-one (Theorem 1), when agents are located on a line, or when agents are split in teams of two equal size (Theorem 4), to cite a few examples. On the contrary, the problem is tractable when the graph is very dense (Theorem 2). It is also easy to see that the problem can be solved efficiently on a *star* network: certainly the center node has to receive her preferred object, and then the remaining question (whether the other agents can each be assigned an object they prefer to the one of the center) turns out to be a matching problem. This gives the intuition that a relevant parameter to study is the size of a *vertex cover* (a subset of nodes in the network including at least one extremity of each edge). For instance the center of a star is a vertex cover. Since at least one of the extremities of each edge is contained in a vertex cover, the rest of the vertices forms a set of pairwise non-adjacent vertices, and thus envy cannot occur within the corresponding set of agents. We provide an algorithm which shows that DEC-LEF is in **XP** (parameterized by the size of a vertex cover) and a proof of **W[1]**-hardness (Theorems 5 and 6). Our findings for DEC-LEF are summarized in Table 1.

degree of G	$\Delta(G) = k$ ($k \geq 1$ fixed)	**NP**-c	Cor. 1
	$\delta(G) = n - k$ ($k \geq 3$ fixed)	**NP**-c	Cor. 2
	$\delta(G) = n - 2$	**P**	Th. 2
number of clusters c in a cluster graph	$c \in \{k, n/k\}$ ($k \geq 2$ fixed)	**NP**-c	Cor. 3
	$c \in \{1, n\}$	**P**	
parameter k on the vertex cover size		**XP**	Th. 5
		W[1]-hard	Th. 6

Table 1: The complexity of DEC-LEF with respect to the degree of its nodes. k is a positive integer and n denotes the number of agents. **P** means polynomial time solvable and **NP**-c means **NP**-complete. $\Delta(G)$ and $\delta(G)$ are the maximum and minimum degrees, respectively, of a vertex in G.

Given that locally envy-free allocations may not exist in the first place, and that the associated decision problems can be hard, it is natural to take an optimization perspective. In our ordinal setting, we shall be concerned with the maximization of the number of non-envious agents, and of a metric averaging the degree of (non-)envy in the society, solely based on the ranks of the items that agents possess. Section 4 is dedicated to optimization problems taking these two different perspectives. We provide approximation algorithms for both approaches. In the first case, we elaborate on the fact that, when an independent set of agents can be identified, a simple sequential picking sequence protocol is sufficient to guarantee that the agents in this set will be locally non-envious (Proposition 1). In fact, this connection to the (maximum) independent set can be further exploited to show that no constant approximation can be found for this objective (Proposition 3), unless $\mathbf{P} = \mathbf{NP}$. In the second case, we build on the observation that random matchings are actually likely to give a high degree of non-envy, and exploit derandomization techniques to obtain a polynomial-time approximation algorithm (Proposition 5).

A variant of DEC-LEF called DEC-LOCATION-LEF is studied in Section 5. This problem asks if one can decide both the placement of the agents (on a given social network) *and* the object allocation so as to satisfy local envy-freeness. For instance, in Example 1, it is natural to imagine that the central authority can also assign agents to their shifts. The problem is (unsurprisingly) shown to be **NP**-complete. A much less expected result, on the positive side, is that the special case of very dense graphs can still be resolved in polynomial time (Theorem 8).

In Section 6 we study the likelihood, for randomly chosen instances of our problems, to be positive (*i.e.* to accept a locally envy-free allocation) –in particular how does it depend on the density of the graph. We first exhibit an asymptotic result (Proposition 6) showing that this event has negligible probability as soon as the degree of the graph is above a fraction $1/e$ of the number of nodes. We complement this by empirical evidence of instances of moderate size (recall that the underlying problems are **NP**-complete), studying in addition how more likely it becomes when the central authority has the extra flexibility to assign agents on the network, as assumed in the DEC-LOCATION-LEF. These experiments are conducted on graphs of regular degree, which are not necessarily realistic. We thus complement our results by using more realistic graphs distributions and restrictions on agents preferences, such as single-peaked domains. In terms of graph structures, we consider scale-free networks, and graphs whose structure depends on the "similarity" between agents (either because similar agents are more likely to be connected, or the other way around).

We provide open problems and future directions in Section 7.

2 Our model and problems

A set of objects O and a set of agents N are given. We assume that $|O| = |N| = n$. This hypothesis on the number of objects being equal to the number of agent does not preclude the case where there are more agents than objects. Indeed, one can add dummy items which will be allocated to the agents who do not receive an object. Each agent i has a preference relation $\succ_i$ over O (a linear order). Let $\succ = (\succ_1, \dots, \succ_n)$ denote the preference profile of the agents. For any positive integer k, $[k]$ stands for $\{1, 2, \dots, k\}$.

We are also given a network modeled as an undirected graph G with vertex set N and edge set E. Each edge in E represents a relation between the corresponding agents. Two agents are directly connected in the network if they can perceive each other and they may envy each other. An instance of a resource allocation problem is thus described by a tuple $\langle N, O, \succ, G = (N, E)\rangle$.

When the network G is dense, it may be easier to describe it through its complement graph $\overline{G}$ which is the unique graph defined on the same vertex set and such that two vertices are connected if and only if they are not connected in G.

The degree of a vertex $v \in N$, denoted by $deg_G(v)$, is the number of edges incident to v. The maximum (respectively, minimum) degree of a graph G, denoted by $\Delta(G)$ (respectively, $\delta(G)$) is the maximum (respectively, minimum) degree of its vertices. A *regular* graph is such that all of its nodes have the same degree. In other words, G is a regular graph if and only if $\delta(G) = \Delta(G)$.

A *partial allocation* $\mathcal{A}$ is a subset of $N \times O$ in which no agent nor object appears twice. If each object and each agent appears exactly once, this partial allocation is called an *allocation*. If agent i appears in $\mathcal{A}$, by an abuse of notation, $\mathcal{A}(i)$ will refer to the object owned by i.

Definition 1 (Locally envy-free) An allocation $\mathcal{A}$ is *locally envy-free* (LEF) if no pair of agents $\{i, j\} \in E$ satisfies $\mathcal{A}(j) \succ_i \mathcal{A}(i)$.

Note that the classical notion of envy-freeness corresponds to the local envy-freeness when graph G is complete. Therefore, the notion of local envy-freeness generalizes the notion of envy-freeness. For a given allocation, an agent is locally envy-free (LEF) if she prefers her object to the object(s) of her neighbor(s).

Several notions of degrees of envy[1] have been studied [21, 23, 26, 41, 42]. In our context we shall study the number of envious agents, and a degree measure capturing some simple notion of intensity of envy, in terms of the difference of ranks[2] between items (these two notions would correspond to $e^{sum,max,bool}$ and $e^{sum,sum,raw}$, up to normalization, under the classification of Chevaleyre et al. [23]).

Definition 2 (Degrees of (non)-envy) Given an allocation $\mathcal{A}$, the degree of envy of agent i towards agent j is

$$e(\mathcal{A}, i, j) = \frac{1}{n-1} \max(0, r_i(\mathcal{A}(i)) - r_i(\mathcal{A}(j)))$$

where $r_i(o)$ is the rank of object o in $i's$ preferences, and $0 \leq e(\mathcal{A}, i, j) \leq 1$. Note that for a given allocation $\mathcal{A}$, an agent i envies a neighboring agent j if and only if $e(\mathcal{A}, i, j) > 0$. Observe also that $e(\mathcal{A}, i, j) = 1$ when i holds an object she ranks last, while j holds an object i ranks first.

[1] The degree of a vertex in a graph should not be confused with the degree of envy which measures how much an agent envies the share of another agent.

[2] This is similar to assuming Borda utilities for the preferences of agents.

Definition 3 (Average degree of (non-)envy) The average degree of envy in the group is

$$\mathcal{E}(\mathcal{A}) = \frac{1}{2\,|E|} \sum_{\{i,j\} \in E} e(\mathcal{A}, i, j) + e(\mathcal{A}, j, i)$$

Respectively, the average degree of non-envy in the group is

$$\mathcal{NE}(\mathcal{A}) = 1 - \mathcal{E}(\mathcal{A}).$$

In other words, we simply average over all the pairs of agents that are connected in the graph (note that the notion of envy being directed, while the underlying graph is not, both directions need to be considered for each edge).

We mainly address four problems: DEC-LEF, MAX-LEF, MAX-NE and DEC-LOCATION-LEF. The first one is a decision problem regarding the existence of an LEF allocation over a given social network. The second and the third ones are optimization problems in which an allocation that is as close as possible to local envy-freeness is sought, using the aforementioned criteria.

Definition 4 (dec-LEF) Given an instance $\langle N, O, \succ, G = (N, E)\rangle$, is there an LEF allocation?

Definition 5 (max-LEF) Given an instance $\langle N, O, \succ, G = (N, E)\rangle$, find an allocation that maximizes the number of LEF agents.

Definition 6 (max-NE) Given an instance $\langle N, O, \succ, G = (N, E)\rangle$, find an allocation $\mathcal{A}$ that maximizes the average degree of non-envy, that is $\mathcal{NE}(\mathcal{A})$.

In DEC-LOCATION-LEF, one has to place the agents on the network in addition to allocate objects to them. This placement makes sense if we consider Example 1 where the agents take shifts.

Definition 7 (dec-location-LEF) Given an undirected network (V, E), and $\langle N, O, \succ\rangle$, are there an allocation $\mathcal{A}$ and a bijection $\mathcal{L} : N \to V$ ($\mathcal{L}$ determines the location of the agents on the network) such that $\mathcal{A}(i) \succ_i \mathcal{A}(j)$ for every edge $\{\mathcal{L}(i), \mathcal{L}(j)\} \in E$?

Example 2 As a warm-up, consider 5 agents located on a line, as depicted below. Each agent has a strict ranking over objects (with top items at the top, e.g. $\succ_1$: $a \succ b \succ c \succ d \succ e$).

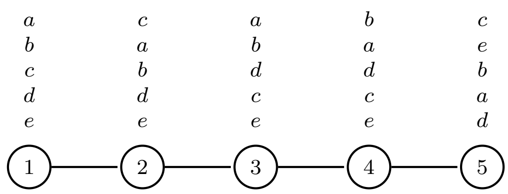

Is there an LEF allocation of goods to agents? If not, what is the minimum number of envious agents? Finally, is it possible to find an LEF allocation by relocating agents on this line?

Let us try to construct an allocation $\mathcal{A}$ that is LEF. Observe that agents 3 and 4, who are neighbors, both rank objects a and b as their first two preferred objects and rank the remaining objects in the last positions of their preference ranking following the same order. This implies that they cannot obtain one of the remaining objects in an LEF allocation, i.e., an object within $\{c, d, e\}$. Indeed, if only one agent between agents 3 and 4 obtains an object in this subset, then she will be envious of the other agent. Otherwise, if they both get an object from this subset, since their preferences over these objects are the same, one of them will necessarily envy the other. Therefore, we have to assign objects a and b to agents 3 and 4 in $\mathcal{A}$, respectively. Consequently, agent 2, neighbor of agent 3, must obtain an object preferred to object a, which is assigned to agent 3. The only object that agent 2 prefers to object a is object c, so we have to assign object c to agent 2 in $\mathcal{A}$. Agent 5, neighbor of agent 4, must get an object preferred to object b, which is assigned to agent 4. The only possible objects are objects c and e, but object c is already assigned to agent 2, thus we assign object e to agent 5 in $\mathcal{A}$. Finally, there only remains object d and agent 1. Agent 1 prefers object c, the object assigned to her neighbor (agent 2), to object d. Therefore, by assigning object d to agent 1 in $\mathcal{A}$, we get that agent 1 is envious of agent 2. Thus, there is no LEF allocation in this instance, implying that this is a *no*-instance of DEC-LEF.

Observe that allocation $\mathcal{A}$ is almost LEF since only agent 1 is envious in $\mathcal{A}$. Therefore, there exists an allocation with only one envious agent. Because there is no LEF allocation, this is the minimum number of envious agents that we can obtain in any allocation. Now in terms of degree of envy, as agent 1 (who holds an object she ranks 4th) only envies agent 2 (who holds agent 1 ranks 3rd), we get that $e(\mathcal{A}, 1, 2) = (4-3)/4 = 1/4$. As this is the only strictly positive envy between any pair of agents, and as there are 4 edges in the network, the average degree of envy is $1/8 \times 1/4 = 1/32$, and the degree of non-envy is thus $31/32$.

Finally, remark that, in allocation $\mathcal{A}$, the only envious agent 1 gets object d, and the only object that agent 1 likes less than d is object e. Object e is owned by agent 5 who is located at a leaf of the path and who, on the opposite, prefers object e to object d. Therefore, by considering a new location of the agents on a path which is the same as the current graph except that agent 1 is a leaf of the path who is connected to agent 5 (i.e., the new path $[2, 3, 4, 5, 1]$), allocation $\mathcal{A} = \{(1, d), (2, c), (3, a), (4, b), (5, e)\}$ is LEF. Hence, this instance is a *yes*-instance of DEC-LOCATION-LEF.

3 Decision problem

This section is devoted to DEC-LEF. Our main findings settle the computational status of DEC-LEF with respect to the degree of the nodes in the network, as well as the size of a vertex cover.

First of all, note that some objects cannot be assigned to certain agents for the allocation to be LEF. For example, the best object of an agent cannot be assigned to one of her neighbors. More generally, no better object than the one allocated to an agent can be assigned to one of her neighbors, leading to the following observations:

Observation 1. *In any LEF allocation, an agent with k neighbors must get an object ranked among her $n-k$ top objects.*

Observation 2. *In any LEF allocation, the best object for an agent is either assigned to herself or to one of her neighbors in $\overline{G}$.*

Observation 1 implies that an agent having $n-1$ neighbors must receive her best object in any LEF allocation. Similarly, agents who do not have any neighbor can receive any object in an LEF allocation.

3.1 DEC-LEF and degree of nodes

Our first result shows that DEC-LEF is computationally difficult, even if the network is very sparse, i.e. each agent has only one neighbor in G (a graph whose every vertex has degree one is called a *matching*). This is somewhat surprising as such a network offers very little possibility for an agent to be envious.

Theorem 1 DEC-LEF *is* **NP**-*complete, even if G is a matching.*

Proof. The reduction is from 3SAT [32]. We are given a set of clauses $C = \{c_1, \cdots, c_m\}$ defined over a set of variables $X = \{x_1, \cdots, x_p\}$. Each clause is disjunctive and consists of 3 literals. The question is whether there exists a truth assignment of the variables which satisfies all the clauses.

Take an instance $\mathcal{I} = \langle C, X \rangle$ of 3SAT and create an instance $\mathcal{J}$ of DEC-LEF as follows.

The set of objects is $O = \{u_i^j : 1 \leq i \leq p,\ 1 \leq j \leq m\} \cup \{\overline{u}_i^j : 1 \leq i \leq p,\ 1 \leq j \leq m\} \cup \{q_j : 1 \leq j \leq m\} \cup \{t_i^j : 1 \leq i \leq p,\ 1 \leq j \leq m\} \cup \{h_\ell : 1 \leq \ell \leq m(p-1)\}$. Here, u_i^j and $\overline{u}_i^j$ correspond to the unnegated and negated literals of x_i possibly present in clause c_j, respectively, q_j corresponds to clause c_j, and the t_i^j's and h_ℓ's are gadgets. Thus, $|O| = 4mp$.

The set of agents N is built as follows. For each $(i,j) \in [p] \times [m]$, create a pair of variable-agents X_i^j, Y_i^j which are linked in the network. For each $j \in [m]$, create a pair of clause-agents K_j, K'_j which are linked in the network. For each $\ell \in [m(p-1)]$, create a pair of garbage-agents L_ℓ, L'_ℓ which are linked in the network. Thus, the network consists of a perfect matching with $4mp$ agents.

Each clause c_j is associated with the pair of clause-agents (K_j, K'_j), q_j and 3 objects corresponding to its literals. For example, $c_2 = x_1 \vee x_4 \vee \overline{x}_5$ is associated with objects q_2, u_1^2, u_4^2, and $\overline{u}_5^2$. The preferences of the clause-agents are:

- $K_j : q_j \succ \ell(j,1) \succ \ell(j,2) \succ \ell(j,3) \succ \text{rest}$
- $K'_j : \ell(j,1) \succ \ell(j,2) \succ \ell(j,3) \succ q_j \succ \text{rest}$

where $\ell(j,i)$ is the object related to the i^{th} literal of c_j, and "rest" means the remaining objects which are arbitrarily ordered, but in the same way for K_j and K'_j.

Each variable x_i is associated with the m pairs of variable-agents (X_i^j, Y_i^j), $1 \leq j \leq m$. The preferences of these variable-agents are:

- $X_i^1 : u_i^1 \succ t_i^1 \succ \overline{u}_i^1 \succ t_i^2 \succ \text{rest}_i^1$
- $Y_i^1 : t_i^1 \succ u_i^1 \succ t_i^2 \succ \overline{u}_i^1 \succ \text{rest}_i^1$

- $X_i^2 : u_i^2 \succ t_i^2 \succ \overline{u}_i^2 \succ t_i^3 \succ \text{rest}_i^2$
- $Y_i^2 : t_i^2 \succ u_i^2 \succ t_i^3 \succ \overline{u}_i^2 \succ \text{rest}_i^2$

- $X_i^3 : u_i^3 \succ t_i^3 \succ \overline{u}_i^3 \succ t_i^4 \succ \text{rest}_i^3$
- $Y_i^3 : t_i^3 \succ u_i^3 \succ t_i^4 \succ \overline{u}_i^3 \succ \text{rest}_i^3$

$\vdots$

- $X_i^{m-1} : u_i^{m-1} \succ t_i^{m-1} \succ \overline{u}_i^{m-1} \succ t_i^m \succ \text{rest}_i^{m-1}$
- $Y_i^{m-1} : t_i^{m-1} \succ u_i^{m-1} \succ t_i^m \succ \overline{u}_i^{m-1} \succ \text{rest}_i^{m-1}$

- $X_i^m : u_i^m \succ t_i^m \succ \overline{u}_i^m \succ t_i^1 \succ \text{rest}_i^m$
- $Y_i^m : t_i^m \succ u_i^m \succ t_i^1 \succ \overline{u}_i^m \succ \text{rest}_i^m$

where "rest_i^j" means the remaining objects arbitrarily ordered, but in the same way for X_i^j and Y_i^j. The preferences of the garbage-agents (L_ℓ, L'_ℓ), $1 \leq \ell \leq m(p-1)$ are:

- $L_\ell : h_\ell \succ U \succ \text{rest}$
- $L'_\ell : U \succ h_\ell \succ \text{rest}$

where $U = \{u_i^j, \overline{u}_i^j : i \in [p],\ j \in [m]\}$, "rest" is the set of remaining objects, and both U and "rest" are arbitrarily ordered in the same way for L_ℓ and L'_ℓ.

Figure 3 summarizes the construction.

We claim that there is an LEF allocation in $\mathcal{J}$ if, and only if, there is a truth assignment satisfying $\mathcal{I}$.

Take a truth assignment which satisfies $\mathcal{I}$. One can allocate objects to each variable-agent pair (X_i^j, Y_i^j) in such a way that it is LEF: If $x_i = true$, then X_i^j gets $\overline{u}_i^j$ and Y_i^j gets t_i^{j+1} (where $t_i^{m+1} := t_i^1$), otherwise $x_i = false$, X_i^j gets u_i^j and Y_i^j gets t_i^j. One can allocate objects to each clause-agent pair (K_j, K'_j)

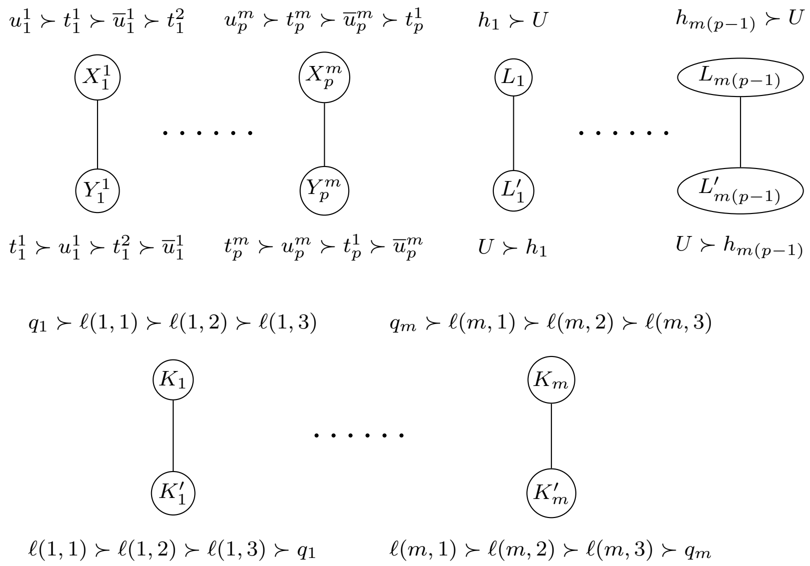

Fig. 3: An overview of the instance of DEC-LEF. Only the most preferred objects of each agent are represented, and the order over the remaining objects is the same for two connected agents. For each $j \in \{1, \dots, m\}$ and $i \in \{1, 2, 3\}$, $\ell(i, j)$ is the object related to the i^{th} literal of c_j, and U represents an arbitrary order over $\{u_i^j, \overline{u}_i^j : i \in [p],\ j \in [m]\}$

in such a way that it is LEF: c_j is satisfied thanks to one of its literals; K_j gets q_j and K'_j gets an unallocated object corresponding to a literal of c_j. Finally, allocate objects to each garbage-agent pair (L_ℓ, L'_ℓ) in such a way that it is LEF: L_ℓ gets h_ℓ and L'_ℓ gets any unallocated objects of U.

Suppose an LEF allocation exists for $\mathcal{J}$. Consider a variable x_i. By construction of the preferences of the variable-agent pair (X_i^1, Y_i^1), we observe that there is absence of envy in only two cases: either (i) X_i^1 gets u_i^1 and Y_i^1 gets t_i^1, or (ii) X_i^1 gets $\overline{u}_i^1$ and Y_i^1 gets t_i^2. If we are in case (i), then there is absence of envy between X_i^m and Y_i^m only if X_i^m gets u_i^m and Y_i^m gets t_i^m because t_i^1 is already allocated, there is absence of envy between X_i^{m-1} and Y_i^{m-1} only if X_i^{m-1} gets u_i^{m-1} and Y_i^{m-1} gets t_i^{m-1} because t_i^m is already allocated, and so on; the X_i^j's get all the u_i^j's (i is fixed but $1 \leq j \leq m$). If we are in case (ii), then there is absence of envy between X_i^2 and Y_i^2 only if X_i^2 gets $\overline{u}_i^2$ and Y_i^2 gets t_i^3 because t_i^2 is already allocated, there is absence of envy between X_i^3 and Y_i^3 only if X_i^3 gets $\overline{u}_i^3$ and Y_i^3 gets t_i^4 because t_i^3 is already allocated, and so on; the X_i^j's get all the $\overline{u}_i^j$'s (i is fixed but $1 \leq j \leq m$). Thus, set x_i to $false$ (respectively, x_i to $true$) if every X_i^j gets u_i^j (respectively, X_i^j gets $\overline{u}_i^j$).

Consider any clause c_j. By construction of the preferences of the clause-agent pair (K_j, K'_j), we observe that there is absence of envy in only three cases: K_j gets q_j and K'_j gets one of the 3 objects associated with the literals of c_j. Since the allocation is LEF, there is some i^* such that K'_j gets either $u^j_{i^*}$ or $\overline{u}^j_{i^*}$.This means that variable x_{i^*} is set to truth if K'_j gets $u^j_{i^*}$ since $u^j_{i^*}$ is not allocated to agent $X^j_{i^*}$, and variable x_{i^*} is set to false if K'_j gets $\overline{u}^j_{i^*}$ since $\overline{u}^j_{i^*}$ is not allocated to agent $X^j_{i^*}$. In both cases, this implies that c_j is satisfied since one of its literal (either x_{i^*} or $\neg x_{i^*}$) is satisfied. To conclude, all the clauses are satisfied. □

The strength of this result lies on the fact that the network structure is extremely simple. As a consequence, it can easily be used as a building block to show hardness of a large variety of graphs. The following lemma shows that one can add a pair of agent/object in an instance of DEC-LEF without changing the complexity of the problem.

Lemma 1 *One can add a pair of agent/object to an instance of* DEC-LEF *without changing the set of LEF allocations (where the additional object is assigned to the additional agent in all solutions). Furthermore, this result does not depend on the set of agents connected to the additional agent in the network (under the condition that she has at least one neighbor).*

Proof. Let $\mathcal{I} = \langle N, O, \succ, G = (N, E)\rangle$ denote an instance of DEC-LEF, and let a and o be the additional agent and object, respectively. The new instance, including a and o, is denoted $\mathcal{J} = \langle N \cup \{a\}, O \cup \{o\}, \succ', G' = (N \cup \{a\}, E')\rangle$, where $\succ' = (\succ'_i)_{i \in N \cup \{a\}}$ are the new preferences of the agents over $O \cup \{o\}$, and E' is the new set of edges of the network. Set E' contains E and does not add a new edge between two agents of N. Furthermore, the set of edges containing a in E' is arbitrary, but contains at least one edge. Let v denote one of the neighbors of agent a in G'. Preference $\succ'_a$ of agent a will be defined as a copy of the preference $\succ_v$ of agent v but with object o at its top. On the contrary, the preference $\succ'_v$ of agent v will be defined as a copy of $\succ_v$ but with object o at its bottom. In other words, the preferences of agent a and v differ only on the position of object o. Finally, preference $\succ'_i$ of any other agent i of N is defined as a copy of $\succ_i$ but with object o at its bottom.

Note first that one can extend any LEF allocation $\mathcal{A}$ of $\mathcal{I}$ into an allocation $\mathcal{A}'$ of $\mathcal{J}$ where each agent of N receives the same object as in $\mathcal{A}$, and where agent a receives object o. Allocation $\mathcal{A}'$ is obviously LEF in $\mathcal{J}$ since agent a receives her most preferred object which is also the last preferred object of any other agent.

We show now that each LEF allocation $\mathcal{A}'$ of $\mathcal{J}$ corresponds to an LEF allocation $\mathcal{A}$ of $\mathcal{I}$. First of all, note that object o should be assigned to agent a in $\mathcal{A}'$. Otherwise, either agent v receives o in $\mathcal{A}'$ and agent a envies agent v, or both agents receive an object of O in $\mathcal{A}'$. In the latter case, one of the agents necessarily envies the other one since they have the same preferences over O, leading to a contradiction with $\mathcal{A}'$ being LEF. Since agent a receives object o in $\mathcal{A}'$, one can easily construct allocation $\mathcal{A}$ by assigning to each agent of N the same object as in $\mathcal{A}'$. Allocation $\mathcal{A}$ is obviously LEF in $\mathcal{I}$ since otherwise $\mathcal{A}'$ would not be LEF in $\mathcal{J}$.

This concludes the proof since we have shown that there is a one-to-one correspondence between the LEF allocations of $\mathcal{I}$ and $\mathcal{J}$. □

As a consequence of Theorem 1 and Lemma 1, the following results hold:

Corollary 1 DEC-LEF *is* **NP***-complete on a line, or on a circle, and generally on graphs of maximum degree k for $k \geq 1$ constant.*

Proof. We only provide a formal proof for the case of the line. The other proofs are similar. We reduce an instance $\mathcal{I}$ of DEC-LEF where the graph is a matching (see Theorem 1) into an instance $\mathcal{J}$ where the graph is a line. Let (v_i, v_i') denote the i^{th} pair or connected agents in the network of instance $\mathcal{I}$, where the order over pairs is arbitrary. Instance $\mathcal{J}$ will be a copy of $\mathcal{I}$ with an additional agent v_i'' for each $i \in \{1, \frac{n}{2} - 1\}$, who will be connected to the agents v_i' and v_{i+1} in the network of $\mathcal{J}$. The network of $\mathcal{J}$ forms a line. The size of $\mathcal{J}$ is at most twice the size of $\mathcal{I}$. According to Lemma 1, one can define the preferences such that the set of LEF allocations of instance $\mathcal{I}$ is the same as the set of LEF allocations of instance $\mathcal{J}$ (except for the additional agents who receive the same additional objects in any LEF allocation). Therefore, the complexity of DEC-LEF is equivalent in instances $\mathcal{I}$ and $\mathcal{J}$. (For the case of a circle, we add an extra agent who will connect the first and last agent of $\mathcal{I}$; and for the case of a graph of maximum degree k, we further connect additional agents to the agents v_i''). □

Given this result, one may suspect the problem to be hard on any graph structure beyond a clique. Our next result shows that if the network is dense enough, then DEC-LEF is polynomial.

Theorem 2 DEC-LEF *in graphs of minimum degree $n-2$ is solvable in polynomial time.*

Proof. Nodes have degree either $n-2$ or $n-1$ in G. In $\overline{G}$, which is the complement graph of G, nodes have degree either 1 or 0. Let $\phi : N \to N$ be such that $\phi(i)$ is the neighbor of i in $\overline{G}$ if i has degree 1 in $\overline{G}$, otherwise $\phi(i) = i$.

We reduce the problem to 2-SAT which is solvable in linear time [4]. Let us consider Boolean variables x_{ij} for $1 \leq i, j \leq n$, such that x_{ij} is true if and only if object j is assigned to i. Denote by o_i^j the object at position j in the preference relation of agent i.

Consider the following formula φ:

$$\varphi \equiv \bigwedge_{i \in N} (x_{io_i^1} \vee x_{io_i^2}) \wedge \bigwedge_{\substack{1 \leq i < \ell \leq n \\ 1 \leq j \leq n}} (\neg x_{ij} \vee \neg x_{\ell j}) \wedge \bigwedge_{i \in N} (x_{io_i^1} \vee x_{\phi(i)o_i^1})$$

The first part of formula φ expresses that each agent must obtain an object within her top 2, as noted in Observation 1. By combination with the second part of φ, we get that the solution must be an assignment: each agent must obtain her first or second choice but not both since every object is owned by at most one agent and $|N| = |O|$. Observation 2 implies that the best object for agent i must be assigned either to agent i or $\phi(i)$. This condition is given by the last part of the formula. Hence, formula φ exactly translates the constraints of an LEF allocation. □

Interestingly, the status of DEC-LEF changes between networks of degree at least $n-2$ and those of degree $n-3$.

Theorem 3 DEC-LEF *is* **NP**-*complete in regular graphs of degree* $n-3$.

Proof. The reduction is from (3, B2)-SAT [12], which is a restriction of 3SAT where each literal appears exactly twice in the clauses, and therefore, each variable appears four times. Take an instance $\mathcal{I} = \langle C, X\rangle$ of (3,B2)-SAT, where $C = \{c_1, \ldots, c_m\}$ is a set of clauses defined over a set of variables $X = \{x_1, \ldots, x_p\}$, and create an instance $\mathcal{J}$ of DEC-LEF as follows.

Instead of describing the network in $\mathcal{J}$, we describe its complementary $\overline{G}$. Note that $\overline{G}$ is a regular graph of degree 2. Hence, $\overline{G}$ contains a collection of cycles. For each variable x_i, we introduce:

- dummy variable-objects q_i^1 and q_i^2,
- literal-objects $u_i^1, u_i^2, \overline{u}_i^1$ and $\overline{u}_i^2$ corresponding to its first and second occurrence as an unnegated and negated literal, respectively,
- a cycle in $\overline{G}$ containing literal-agents $X_i^1, \overline{X}_i^1, X_i^2$ and $\overline{X}_i^2$, connected in this order.

We denote by X_i the subset of literal-agents containing $X_i^1, \overline{X}_i^1, X_i^2$ and $\overline{X}_i^2$. The preferences of the literal-agents are as follows, for each $i \in [p]$:

- $X_i^1 : q_i^1 \succ q_i^2 \succ u_i^1 \succ \ldots$
- $\overline{X}_i^1 : q_i^1 \succ q_i^2 \succ \overline{u}_i^1 \succ \ldots$
- $\overline{X}_i^2 : q_i^2 \succ q_i^1 \succ \overline{u}_i^2 \succ \ldots$
- $X_i^2 : q_i^2 \succ q_i^1 \succ u_i^2 \succ \ldots$

Note that only the 3 top objects are represented since no object ranked below can lead to an LEF allocation (see Observation 1). We show that in any LEF allocation, either q_i^1 and q_i^2 are allocated to agents X_i^1 and X_i^2, respectively, or q_i^1 and q_i^2 are allocated to agents $\overline{X}_i^1$ and $\overline{X}_i^2$, respectively. For any $j \in \{1,2\}$, if q_i^j is allocated to agent $Y \notin X_i$, then agents X_i^j will envy agent Y because they are neighbors in G and q_i^j is the most favorite object of agent X_i^j. Moreover, if q_i^j is owned by agent X_i^{3-j} (respectively, $\overline{X}_i^{3-j}$) then agent X_i^j (respectively, $\overline{X}_i^j$) will be envious of agent X_i^{3-j} (respectively, $\overline{X}_i^{3-j}$) because they are neighbors in G and q_i^j is the most favorite object of agent X_i^j (respectively, $\overline{X}_i^j$). Therefore, q_i^j is assigned to either agent X_i^j or $\overline{X}_i^j$. Finally, if agent X_i^j (respectively, $\overline{X}_i^j$) receives q_i^j and agent X_i^{3-j} (respectively, $\overline{X}_i^{3-j}$) does not receive q_i^{3-j} then agent X_i^{3-j} (respectively, $\overline{X}_i^{3-j}$) will envy agent

X_i^j since they are neighbors in G and agent X_i^{3-j} (respectively, $\overline{X}_i^{3-j}$) did not receive her most favorite object and her second most favorite object is q_i^j.

The case where q_i^1 and q_i^2 are allocated to agents X_i^1 and X_i^2, respectively, can be interpreted in $\mathcal{I}$ as setting x_i to true, and the case where q_i^1 and q_i^2 are allocated to agents $\overline{X}_i^1$ and $\overline{X}_i^2$, respectively, as setting x_i to false.

For each clause c_j we introduce:

- dummy clause-objects d_j^1 and d_j^2,
- a cycle in $\overline{G}$ containing clause-agents K_j^1, K_j^2, and K_j^3.

The preferences of clause-agent K_j^i, for $j \in [m]$ and $i \in [3]$, are:

- $K_j^i : d_j^1 \succ d_j^2 \succ \ell(j,i) \succ \dots$

where $\ell(j,i)$ is the literal-object corresponding to the i^{th} literal of c_j. We denote by K_j the subset of clause-agents containing K_j^1, K_j^2 and K_j^3. We show that an allocation is LEF if d_j^1, d_j^2 and one literal-object corresponding to a literal of c_j are assigned to the agents of K_j. For any $i \in \{1,2\}$, if d_j^i is allocated to agent $Y \notin K_j$ then one of the agents of K_j receives neither d_j^1 nor d_j^2 and will envy agent Y who is her neighbor in G. Therefore, objects d_j^1 and d_j^2 are assigned among the agents of K_i. If the agent of K_i who receives neither d_j^1 nor d_j^2, say K_j^i, does not receive $\ell(i,j)$ then she will envy the agent who receives $\ell(i,j)$ and who necessarily is her neighbor in G. This gadget can be interpreted in $\mathcal{I}$ as the requirement for at least one literal of c_j to be true.

Figure 4 summarizes the agents of the reduction introduced so far.

The reduction is almost complete but it remains to describe gadgets collecting all unassigned objects. Indeed, so far we have introduced $4p+3m$ agents and $6p+2m$ objects. It remains to construct garbage collectors for the $2p-m$ remaining objects. Note that $2p-m \geq 0$ holds since each variable appears in 4 clauses and each clause contains 3 literals (in other words, $4p = 3m$ holds). Note also that no dummy object (neither variable nor clause) may be part of the remaining objects since they must be assigned to literal-agents or clause-agents in any LEF allocation. Let $\mathcal{L} = \{u_i^j, \overline{u}_i^j : i \in [p], j \in [2]\}$ denote the set of literal-objects, where literal-objects are ordered arbitrarily, and let $\mathcal{L}(i)$ denote the i^{th} element of $\mathcal{L}$.

Let us now describe a gadget collecting a single object of $\mathcal{L}$. For each $i \in [4p]$, we introduce:

- objects t_i^1 and t_i^2,
- a cycle in $\overline{G}$ containing gadget-agents L_i^1, L_i^2 and L_i^3.

Furthermore, for each $i \in [4p-1]$, we introduce gadget-object h_i. Globally, in this gadget, we introduce $12p$ new agents and $12p-1$ new objects. Preferences are as follows, for each $i \in [4p]$ (where h_0 and h_{4p} stand for h_1 and h_{4p-1}, respectively):

- $L_i^1 : t_i^1 \succ t_i^2 \succ h_{i-1} \succ \dots$
- $L_i^2 : t_i^1 \succ t_i^2 \succ \mathcal{L}(i) \succ \dots$

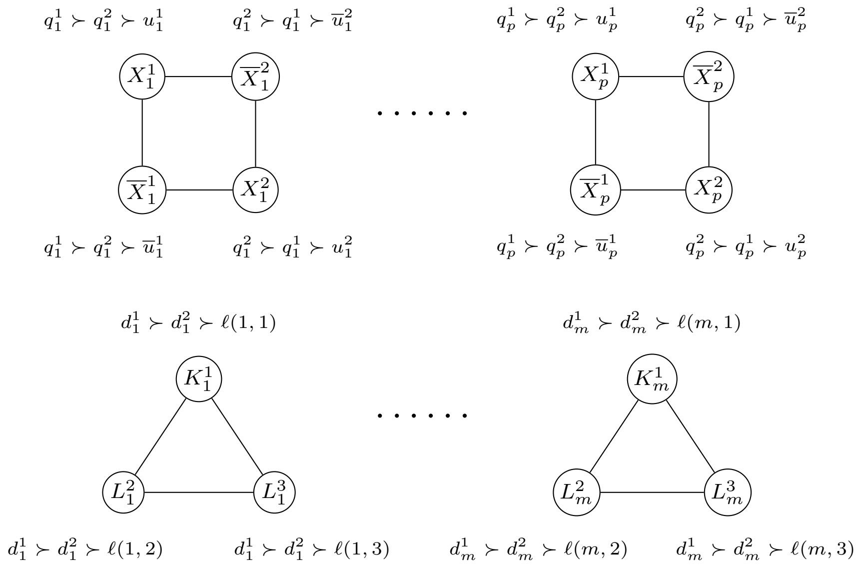

Fig. 4: A partial description of the graph of non-envy $\overline{G}$. Note that the neighborhood of each agent in G corresponds to the whole set of agents except for her two neighbors in $\overline{G}$ (described in this figure). Only the most preferred objects of each agent are represented. For each $j \in \{1, \dots, m\}$ and $i \in \{1, 2, 3\}$, $\ell(i, j)$ is the literal-object corresponding to the i^{th} literal of c_j.

– $L_i^3 : t_i^1 \succ t_i^2 \succ h_i \succ \dots$

Note that in any LEF allocation, objects t_i^1 and t_i^2 are allocated to agents belonging to $\{L_i^1, L_i^2, L_i^3\}$, and the remaining unassigned agent receives either h_{i-1}, h_i or $\mathcal{L}(i)$ (the proof is similar as the above proof for the clause-agents). Since no more than $4p - 1$ agents can receive a gadget-object, at least one literal-object is assigned to agent L_i^2 for some $i \in [4p]$. Moreover, all gadget-objects must be assigned to gadget-agents since no other agent has a gadget-object in her top 3 objects. Therefore, in every LEF allocation, exactly one literal-object is allocated to an agent belonging to the gadget.

Figure 5 provides a graphical description of this gadget.

Now let us show that one can allocate objects without envy in the gadget. Let $\mathcal{L}(i)$ be the literal-object assigned in the gadget. This object must be assigned to L_i^2. Assign objects t_i^1 and t_i^2 to agents L_i^1 and L_i^3, respectively. For any $j \neq i$, assign object t_j^1 to agent L_j^2. Finally, for any $j > i$, object h_{j-1} is assigned to L_j^1 and t_j^2 is assigned to L_j^3, and for any $j < i$, object h_j is assigned to L_j^3 and object t_j^2 is assigned to L_j^1.

We use exactly $2p - m$ copies of this gadget in order to collect all the remaining literal-objects of the first part of the construction, and thus obtaining as many agents as objects in the whole reduction.

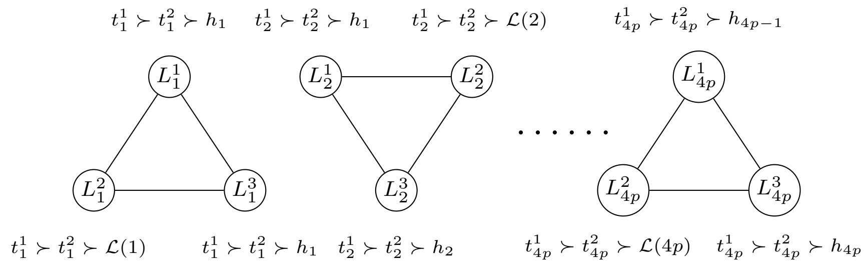

Fig. 5: The graph of non-envy for the gadget aiming to absorb one literal-object not assigned to an agent of Figure 4. For each $i \in \{1, \dots, 4p\}$, $\mathcal{L}(i)$ is the i^{th} literal-object.

We claim that C is satisfiable in instance $\mathcal{I}$ if and only if $\mathcal{J}$ has an LEF allocation.

Suppose first that there exists a truth assignment ϕ of the variables in X which satisfies all clauses in C. For each variable x_i which is true (respectively, false) in ϕ, we assign objects q_i^1 to agent X_i^1 (respectively, $\overline{X}_i^1$), object q_i^2 to agents X_i^2 (respectively, $\overline{X}_i^2$), object $\overline{u}_i^1$ (respectively, u_i^1) to agents $\overline{X}_i^1$ (respectively, X_i^1) and object $\overline{u}_i^2$ (respectively, u_i^2) to agent $\overline{X}_i^2$ (respectively, X_i^2). Note that each agent of X_i receives either her most preferred object, or receive her third most preferred and her two neighbors in $\overline{G}$ receive her first and second most preferred object. Therefore, no agent of X_i envies one of her neighbors in G. Note also that the unassigned literal-objects are associated with literals which are true according to ϕ. Since each clause c_j is satisfied by ϕ, there exists at least one unassigned literal-object that we assign to clause-agents K_j^i, where i is the index of its corresponding literal in c_j. Note that this is her third most preferred object. The two other clause-agents of K_j receive their first and second most preferred objects i.e., dummy-objects d_j^1 and d_j^2. Since the three agents of K_j are not neighbors in G, none of them envies one of her neighbors. Finally, it suffices to assign the remaining literal-objects to garbage-agents, as previously described in the construction of the gadgets, in such a way that no garbage-agent can be envious. We obtain an LEF allocation.

Suppose now that there exists an LEF allocation. As shown above, in any LEF allocation, objects q_i^1 and q_i^2 must be assigned either to (i) agents X_i^1 and X_i^2, respectively, or to (ii) agents $\overline{X}_i^1$ and $\overline{X}_i^2$, respectively. In case (i), literal-objects $\overline{u}_i^1$ and $\overline{u}_i^2$ must be assigned to agents $\overline{X}_i^1$ and $\overline{X}_i^2$, respectively. In case (ii), literal-objects u_i^1 and u_i^2 must be assigned to agents X_i^1 and X_i^2, respectively. Let ϕ denote the truth assignment of the variables of X such that for each variable x_i, if literal-objects are assigned to the variable-agents of X_i as in case (i) then x_i is set to true, and otherwise x_i is set to false. We claim that ϕ satisfies all clauses in C. Indeed, we have shown above that in any LEF allocation, dummy clause-objects d_j^1 and d_j^2, as well as one literal-object corresponding to one literal appearing in c_j, must be assigned to the agents of K_j. This literal-object is true according to ϕ since the corresponding literal-object is not assigned to a variable-agent. Therefore, ϕ satifies clause c_j. □

In the same vein as for Theorem 1, the hardness result of Theorem 3 can be extended to more general classes of graphs.

Corollary 2 DEC-LEF *is* **NP***-complete on graphs of minimum degree* $n-k$ *for* $k \geq 3$ *constant.*

Proof. The proof is similar to the proof of Corollary 1 except that we add only one additional agent who is connected to $n-k$ agents chosen arbitrarily (note that n refers to the number of agents in the new instance). □

Related to the question of the degree of the nodes, it appears interesting to determine how the computational hardness of DEC-LEF evolves on *cluster graphs* (such graphs are collections of disjoints cliques). The cluster graphs are relevant in the context of a social network, because they may represent several groups of agents that do not have interconnections (e.g. families or different sport teams). In fact, the problem is computationally hard when the cluster graph is composed of $n/2$ cliques because this is the case of the matching (Theorem 1). This hardness result is extended to any cluster graph composed of n/k cliques (for $k \geq 2$ constant) according to the construction used to obtain Corollary 1. Note that the case of n clusters is trivial since it is the empty graph. Moreover, the problem is solvable in polynomial time when there is only one clique in the cluster graph (the easy case of the complete graph). A natural question is then the complexity of DEC-LEF when the cluster graph is only composed of two cliques. The next theorem shows that even in this case, the problem is **NP**-complete.

Theorem 4 DEC-LEF *is* **NP***-complete even when the social network is restricted to two cliques of equal size.*

Proof. The reduction is from an instance $\mathcal{I}$ of 3SAT. Let $C = \{c_1, \ldots, c_m\}$ and $X = \{x_1, \ldots, x_p\}$ denote the set of clauses and variables, respectively. The reduction to an instance $\mathcal{J}$ of DEC-LEF is as follows. Let Q_1 and Q_2 denote the two cliques of G that we are going to construct. We introduce

- two agents Q_1^1 and Q_1^2 belonging to Q_1,
- two agents Q_2^1 and Q_2^2 belonging to Q_2,
- four objects q_1^1, q_1^2, q_2^1 and q_2^2.

The preferences of agents Q_i^1 and Q_i^2, for each $i \in \{1, 2\}$, are:

- $Q_i^1 : q_i^1 \succ q_i^2 \succ rest$
- $Q_i^2 : q_i^2 \succ q_i^1 \succ rest$

where *rest* is an arbitrary order over the objects different from q_i^1 and q_i^2, but in the same order for both agents Q_i^1 and Q_i^2. We show that in any LEF allocation, agents Q_i^1 and Q_i^2, who are neighbors in G, will receive objects q_i^1 and q_i^2, respectively. Assume that agent Q_i^j, with $i, j \in \{1, 2\}$, receives object $o \neq q_i^j$ in an LEF allocation. If agent Q_i^{3-j} receives object $o' \notin \{q_i^1, q_i^2\}$ then $o' \succ_{Q_i^j} o \Leftrightarrow o' \succ_{Q_i^{3-j}} o$ holds and either Q_i^j envies Q_i^{3-j} or Q_i^{3-j} envies Q_i^j, a contradiction. On the other hand, if agent Q_i^{3-j} receives either q_i^1 or q_i^2 then agent Q_i^j envies Q_i^{3-j}, a contradiction. Assume now that agent Q_i^j, with $i, j \in \{1, 2\}$, receives object q_i^{3-j}. Since q_i^{3-j} is the most preferred object for agent Q_i^{3-j} and Q_i^j is a neighbor of Q_i^{3-j} in G, agent Q_i^{3-j} envies Q_i^j, a contradiction. Finally, if q_i^j receives object q_i^j and agent Q_i^{3-j} does not receive object q_i^{3-j} then agent Q_i^{3-j} envies Q_i^j, a contradiction. Therefore, agent Q_i^j receives object q_i^j for any $i, j \in \{1, 2\}$. As we will see later, the fact that the allocation of each objects q_i^j, with $i, j \in \{1, 2\}$, is fixed will enforce the other agents to obtain a more preferred object according to their preference.

For each variable x_i we introduce:

- variable-agent X_i^1 belonging to clique Q_1,
- variable-agent X_i^2 belonging to clique Q_2,
- variable-objects u_i and $\bar{u}_i$.

The preferences of variable-agent X_i^j, for each $i \in \{1, \ldots, p\}$ and $j \in \{1, 2\}$, are:

- $X_i^j : u_i \succ \bar{u}_i \succ q_j^1 \succ rest$

where *rest* is an arbitrary order over the remaining objects. We show that in any LEF allocation, objects u_i and $\bar{u}_i$ are assigned to agents X_i^j and X_i^{3-j}, respectively, with j either equal to 1 or 2. Assume by contradiction that agent X_i^j, with $i \in \{1, \ldots, q\}$ and $j \in \{1, 2\}$, receives object $o \notin \{u_i, \bar{u}_i\}$ in an LEF allocation. Since the allocation is LEF, agent Q_j^1 must receive object q_j^1 (see the proof above). However, agents X_i^j and Q_j^1 are neighbors in G and agent X_i^j prefers q_j^1 to o, and therefore agent X_i^j envies Q_j^1, a contradiction. The case where u_i and $\bar{u}_i$ are assigned to agents X_i^1 and X_i^2, respectively, can be interpreted in $\mathcal{I}$ as setting variable x_i to true, and the case where u_i and $\bar{u}_i$ are assigned to agents X_i^2 and X_i^1, respectively, can be interpreted in $\mathcal{I}$ as setting variable x_i to false.

For each clause c_j we introduce:

- clause-agents K_j^1, K_j^2 and K_j^3 belonging to Q_1,
- clause-agent K_j and dummy agents L_j^1 and L_j^2 belonging to Q_2,
- clause-objects t_j, t_j^1, t_j^2 and t_j^3,
- dummy objects h_j^1 and h_j^2.

The preferences of clause-agents K_j^i and K_j and dummy agent L_j^ℓ, for each $j \in \{1, \ldots, m\}$, $i \in \{1, 2, 3\}$ and $\ell \in \{1, 2\}$, are:

- $K_j^i : t_j^i \succ \bar{\ell}(j, i) \succ t_j \succ q_1^1 \succ rest$

- $K_j : t_j^1 \succ t_j^2 \succ t_j^3 \succ q_2^1 \succ rest$
- $L_j^\ell : h_j^\ell \succ q_2^1 \succ rest$

where $\bar{\ell}(j,i)$ denotes the variable-object corresponding to the negation of literal i in clause c_j, and $rest$ is an arbitrary order over the remaining objects. It is easy to show that each dummy agent L_j^ℓ should receive object h_j^ℓ for an allocation to be LEF since agent Q_2^1 should receive object q_2^1. For the same reason, agent K_j should receive either t_j^1, t_j^2 or t_j^3. Assume that K_j receives t_j^i. In that case, agent K_j^i should receive item t_j for the allocation to be LEF since $\bar{\ell}(j,\ell)$ should be assigned to a variable-agent and q_1^1 should be assigned to agent Q_1^1 who is a neighbor of K_j^i in G. Furthermore, for the allocation to be LEF, $\bar{\ell}(j,i)$ should not be assigned to an agent of Q_j since otherwise, K_j^i would be envious of this agent. This gadget can be interpreted in $\mathcal{I}$ as the requirement for at least one literal of c_j to be true.

agents of Q_1	preferences	agents of Q_2	preferences
Q_1^1	$q_1^1 \succ q_1^2$	Q_2^1	$q_2^1 \succ q_2^2$
Q_1^2	$q_1^2 \succ q_1^1$	Q_2^2	$q_2^2 \succ q_2^1$
X_1^1	$u_1 \succ \bar{u}_1 \succ q_1^1$	X_1^2	$u_1 \succ \bar{u}_1 \succ q_2^1$
...	...	...	...
X_p^1	$u_p \succ \bar{u}_p \succ q_1^1$	X_p^2	$u_p \succ \bar{u}_p \succ q_2^1$
K_1^1	$t_1^1 \succ \bar{\ell}(1,1) \succ t_1 \succ q_1^1$	K_1	$t_1^1 \succ t_1^2 \succ t_1^3 \succ q_2^1$
...	...	...	...
K_m^1	$t_m^1 \succ \bar{\ell}(m,1) \succ t_m \succ q_1^1$	K_m	$t_m^1 \succ t_m^2 \succ t_m^3 \succ q_2^1$
K_1^2	$t_1^3 \succ \bar{\ell}(1,3) \succ t_1 \succ q_1^1$	L_1^1	$h_1^1 \succ q_2^1$
...	...	...	...
K_m^2	$t_m^3 \succ \bar{\ell}(m,3) \succ t_m \succ q_1^1$	L_m^1	$h_m^1 \succ q_2^1$
K_1^3	$t_1^3 \succ \bar{\ell}(1,3) \succ t_1 \succ q_1^1$	L_1^2	$h_1^2 \succ q_2^1$
...	...	...	...
K_m^3	$t_m^3 \succ \bar{\ell}(m,3) \succ t_m \succ q_1^1$	L_m^2	$h_m^2 \succ q_2^1$

Table 2: The agents of clique Q_1 (clique Q_2, respectively) are listed in the first column (third column, respectively) and their preferences are given in the second column (fourth column, respectively). For any $i \in \{1,\ldots,m\}$ and $j \in \{1,2,3\}$, $\bar{\ell}(j,i)$ denotes the variable-object corresponding to the negation of literal i in clause c_j

We claim that C is satisfiable in $\mathcal{I}$ if and only if $\mathcal{J}$ contains an LEF allocation. Suppose first that there exists truth assignment ϕ of X that satisfies each clause of C. We construct from ϕ an LEF allocation in $\mathcal{J}$. Assign q_i^j to agent Q_i^j for each $i,j \in \{1,2\}$. Furthermore, assign h_j^ℓ to agent L_j^ℓ for each $j \in \{1,\ldots,m\}$ and $\ell \in \{1,2\}$. Since each of these agents receives her most preferred object, none of them will be envious. For each variable x_i, assign u_i and $\bar{u}_i$ to agents X_i^1 and X_i^2, respectively, if x_i is true in ϕ, and otherwise assign u_i and $\bar{u}_i$ to agents X_i^2 and X_i^1, respectively. Since X_i^1 and X_i^2 are not neighbors in G and u_i and $\bar{u}_i$ are their two most favorite objects, neither of them will be envious. Finally, for each clause c_j, pick one literal which is true according to ϕ, say literal i, and assign t_j^i and t_j to agents K_j and K_j^i, respectively. Furthermore, assign to the remaining clause-agents in Q_1 their most favorite objects. It is easy to check that none of these agents will be envious, and the resulting allocation is LEF.

Suppose now that there exists an LEF allocation $\mathcal{A}$ for $\mathcal{J}$. We construct from $\mathcal{A}$ a truth assignment ϕ which satisfies each clause of C. As shown above, in $\mathcal{A}$ either u_i or $\bar{u}_i$ is assigned to agent X_i for each $i \in \{1, \ldots, p\}$. Therefore, set to true in ϕ each variable x_i such that u_i is assigned to X_i, and set to false in ϕ each variable x_i such that $\bar{u}_i$ is assigned to X_i. As shown above, for each $j \in \{1, \ldots, m\}$, there exists agent K_j^i who receives object t_j in $\mathcal{A}$. Since K_j^i is not envious, then no variable-agent receives $\bar{\ell}(j,i)$. Therefore, literal $\ell(j,i)$ is true according to ϕ, and each clause of C is satisfied by ϕ. □

By adding clusters of dummy agents having their associated dummy resource on top of their preference ranking, we can generalize the previous negative result to any cluster graph with $k \geq 2$ (k constant) clusters.

Corollary 3 DEC-LEF *is* **NP**-*complete in any cluster graph with* $k \geq 2$ *clusters or* n/k $(k \geq 2)$ *clusters for* k *constant.*

Proof. As for Lemma 1, we show that we can add an agent, who is isolated inthe network, without changing the set of LEF allocations (under the condition that each vertex of the network had at least one neighbor). Let $\mathcal{I}$ denote the original instance of DEC-LEF, and let $\mathcal{J}$ denote the new instance obtained after adding agent a and object o. The preferences of agent a in $\mathcal{J}$ are arbitrary but object o must be on top. On the other hand, the other agents keep in $\mathcal{J}$ the same preferences as in $\mathcal{I}$ except that object o is at the bottom of their preferences. It is easy to check that there is a one-to-one correspondence between the set of allocations of $\mathcal{I}$ and $\mathcal{J}$. This is mainly due to the fact that no agent, except a, can receive object o without envying one of her neighbors (object o is her least preferred object). Furthermore, every LEF allocation of $\mathcal{I}$ can be completed by allocating object o to agent a without creating envy.

This means that, starting from an instance of DEC-LEF where the network is a clique, on can add one by one isolated cliques of the same size without changing the set of LEF allocation. Indeed, an additional clique, say K, can be added vertex by vertex, starting with an isolated vertex. According to the first paragraph, the first vertex can be added without changing the set of LEF allocations. Furthermore, the remaining vertices will be connected to the vertices of K already present in the graph, and therefore Lemma 1 implies that this can be done without changing the set of LEF allocation. All in all, one can add $k-1$ isolated cliques without changing the set of LEF allocations, leading to an equivalent instance of DEC-LEF containing k cliques of the same size. □

3.2 DEC-LEF and vertex cover

So far the complexity of DEC-LEF has been investigated through the degree of its nodes, but other parameters can be taken into account. Let us show how the size of a (smallest) *vertex cover* can help. A *vertex cover* C of $G = (N, E)$ is a subset of nodes such that $\{u, v\} \cap C \neq \emptyset$ for every edge $\{u, v\} \in E$. Since at least one of the extremities of each edge is contained in C, $I := N \setminus C$ must be an *independent set*, that is a set of pairwise non-adjacent vertices. Thus, an agent of I can only envy an agent of C.

Theorem 5 *If the network G admits a vertex cover of size k, then* DEC-LEF *can be answered in $\mathcal{O}(n^{k+3})$.*

Proof. Let C be a vertex cover of the network and let $I := N \setminus C$ be the corresponding independent set. One can easily construct a vertex cover of size k (if such a set exists) by testing every subset of vertices of size k. The complexity of such a brute force algorithm is in $\mathcal{O}(n^{k+3})$ since there is at most n^k subsets of vertices of size k and testing if each of them is a vertex cover can be done in $\mathcal{O}(n^2)$ by checking if each edge is covered.

Then, use brute force to assign k objects of O to the agents of C with time complexity in $\mathcal{O}(n^k)$. For each partial allocation $\mathcal{A}$ without envy within C, let $O_{-\mathcal{A}}$ be the set of unassigned objects (if no such partial allocation exists, then we can immediately conclude that no LEF allocation exists). Build a bipartite graph $(I, O_{-\mathcal{A}}; E')$ with an edge from agent $i \in I$ to object $o \in O_{-\mathcal{A}}$ if assigning o to i does not create envy. There is an LEF allocation which extends $\mathcal{A}$ if and only if the bipartite graph admits a perfect matching. The existence of a perfect matching in a bipartite graph can be checked in $\mathcal{O}(n^3)$ (see e.g., the book of Burkard, Dell'Amico and Martello [20]). □

The method is efficient when k is small. For instance, DEC-LEF is polynomial if the network is a star because the central node of a star is a vertex cover. More generally, Theorem 5 implies that DEC-LEF is polynomial when $k = \mathcal{O}(1)$.

Theorem 5 implies that DEC-LEF belongs to **XP** when the fixed parameter under consideration is the size of a vertex cover. Recall that a problem belongs to **FPT** if there is an algorithm to solve it with time complexity in $\mathcal{O}(f(k)n^c)$, where c is a constant value and f is an arbitrary function depending only on k. One could expect that DEC-LEF also belongs to **FPT** for the same parameter since the problem of finding a vertex cover of size k is **FPT** [39]. However, the following theorem shows that there is no hope that DEC-LEF belongs to **FPT**.

Theorem 6 DEC-LEF *parameterized by the size of a vertex cover is* **W[1]**-*hard.*

Proof. We present a parameterized reduction from MULTICOLORED INDEPENDENT SET [28]. An instance $\mathcal{I}$ of MULTICOLORED INDEPENDENT SET consists of a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, an integer k, and a partition $(\mathcal{V}_1, \ldots, \mathcal{V}_k)$ of $\mathcal{V}$. The task is to decide if there is an independent set of size k in $\mathcal{G}$ containing exactly one vertex from each set $\mathcal{V}_i$. Let m and p denote the number of vertices and edges in $\mathcal{G}$, respectively.

We construct an instance $\mathcal{J}$ of DEC-LEF as follows. For each vertex v in $\mathcal{V}$, we introduce object o_v. Let O_i denote the set of objects $\{o_v : v \in \mathcal{V}_i\}$, and let $O_i^{\uparrow}$ denote an arbitrary order over the objects of O_i. For each edge $e = \{v, v'\}$ in $\mathcal{E}$, we introduce two agents X_e^v and $X_e^{v'}$, and two edge-objects o_e and o'_e. Let $O_{\mathcal{E}}$ denote the set of edge-objects, and let $O_{\mathcal{E}}^{\uparrow}$ denote an arbitrary ranking over the objects of $O_{\mathcal{E}}$.

For each integer $i \leq k$, we introduce agent K_i. The agents of $\{K_i\}_{i \leq k}$ form a clique in the network G. Furthermore, for each vertex $v \in V_i$ and for each edge $e = \{v, v'\}$ in $\mathcal{E}$, agent X_e^v is connected to agent K_i in G. Finally, for each integer $j \leq |\mathcal{V}| - k$, we introduce agent D_j who is isolated in G. All in all, there are $m + 2p$ agents and objects.

Preferences are the following:

- $K_i : O_i^{\uparrow} \succ O_1^{\uparrow} \succ \ldots \succ O_{i-1}^{\uparrow} \succ O_{i+1}^{\uparrow} \succ \ldots \succ O_k^{\uparrow} \succ O_{\mathcal{E}}^{\uparrow}$
- $X_e^v : o_e \succ o_v \succ o'_e \succ \ldots$

Since agent D_j is isolated in G, her preferences may be arbitrary. It is easy to check that $\{K_i\}_{i \leq k}$ forms a vertex cover in the network.

Figure 6 summarizes the construction.

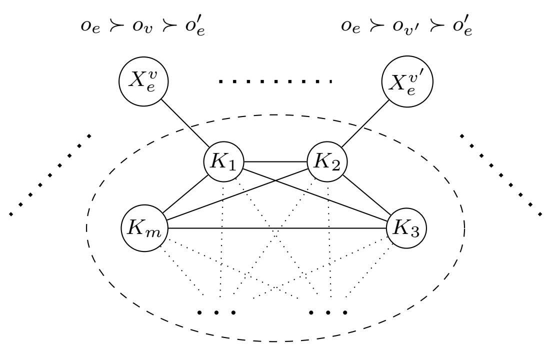

Fig. 6: The graph of envy constructed by the reduction. Agents $K_1, K_2, \ldots, K_m$ form a clique and the agents out of this clique are only connected to one agent belonging to the clique. We assume that $e = (v, v')$ is an edge of $\mathcal{E}$, v belongs to $\mathcal{V}_1$ and v' belongs to $\mathcal{V}_2$.

We show that $\mathcal{G}$ has an independent set of size k containing one vertex in each set $\mathcal{V}_i$ if and only if an LEF allocation exists in $\mathcal{J}$. Assume first that $\{v_1, \ldots, v_k\}$ is an independent set in $\mathcal{G}$, where $v_i \in \mathcal{V}_i$ for each $i \leq k$. We construct an LEF allocation in $\mathcal{J}$ as follows. For each $i \leq k$, assign o_{v_i} to K_i, and for each edge $e = (v_i, v')$ in $\mathcal{E}$, assign o_e to $X_e^{v_i}$. For each agent X_e^v such that v is not selected in the independent set (i.e., $v \notin \{v_1, \ldots, v_k\}$), assign o_e to X_e^v if it is still available, and otherwise assign o'_e to X_e^v. Finally, assign the remaining objects arbitrarily. We claim that this allocation is envy-free. Indeed, each agent K_i receives an object of O_i. Therefore, no agent K_i will envy another agent K_j with $j \neq i$. Furthermore, for each vertex v in $\mathcal{V}_i$ and for each edge e in $\mathcal{E}$, agent X_e^v has a single neighbor who is K_i. If K_i receives o_{v_i} and $v = v_i$ then $X_e^{v_i}$ receives o_e, and otherwise X_e^v receives o'_e. In both cases, agent K_i does not envy X_e^v since X_e^v receives an object of $O_{\mathcal{E}}$, and agent X_e^v does not envy agent K_i since agent X_e^v receives either her most favorite object or her third most favorite object while agent K_i does not receive o_v (the second most favorite object of agent X_e^v).

Assume now that an LEF allocation $\mathcal{A}$ exists in $\mathcal{J}$. We claim that each agent K_i should receive an object of O_i in $\mathcal{A}$. By contradiction, assume that agent K_i receives object $o \notin O_i$ in $\mathcal{A}$. Note that for any $j \neq i$, K_i and K_j are neighbors. Hence, for any object o', if $o \notin O_j$ and $o' \notin O_i \cup O_j$ then $o \succ_{K_i} o'$ if and only if $o \succ_{K_j} o'$ holds. Furthermore, if agent K_j receives an object of O_i in $\mathcal{A}$ then agent K_i will envy her. This implies that if $o \notin O_j$ then an object of O_j must be assigned to K_j in $\mathcal{A}$ to avoid envy between agents K_i and K_j. Therefore, if $o \in O_{\mathcal{E}}$ then each agent K_j, with $j \neq i$, receives an object of O_j in $\mathcal{A}$ and agent K_i will envy them, a contradiction. On the other hand, if $o \in O_j$ for some $j \neq i$ then agent K_j receives an object of O_j in $\mathcal{A}$ and either K_i envies K_j or K_j envies K_i since $o \succ_{K_i} o'$ if and only if $o \succ_{K_j} o'$ holds, a contradiction. Hence, each agent K_i should receive an object of O_i in $\mathcal{A}$. Let o_{v_i} denote the object assigned to K_i in $\mathcal{A}$. We claim that $\{v_1, \ldots, v_k\}$ forms an independent set in $\mathcal{G}$. By contradiction assume that an edge, say e, connects v_i and v_j in $\mathcal{G}$. This implies by construction that $X_e^{v_i}$ and $X_e^{v_j}$ are neighbors of K_i and K_j in G, respectively. On one hand, if $X_e^{v_i}$ does not receive o_e in $\mathcal{A}$ then she envies K_i who receives v_i. On the other hand, if $X_e^{v_j}$ does not receive o_e in $\mathcal{A}$ then she envies K_j who receives v_j. Therefore, o_e must be assigned to both $X_e^{v_i}$ and $X_e^{v_j}$, leading to a contradiction since o_e cannot be assigned twice. □

4 Optimization

In light of Section 3, we know that both MAX-LEF and MAX-NE are **NP**-hard even on very simple graph structures. We present in this section approximation algorithms for MAX-LEF and MAX-NE.

4.1 Maximizing the number of LEF agents

This subsection is dedicated to MAX-LEF, which aims at maximizing the number of non-envious agents. A general method is proposed in Algorithm 1. For a maximization problem, an algorithm is ρ-approximate, with $\rho \in [0, 1]$, if it outputs a solution whose value is at least ρ-times the optimal value, for any instance.

Algorithm 1:

Data: An instance $\langle N, O, \succ, G = (N, E)\rangle$
Result: An allocation $\mathcal{A}$
1 Let $\mathcal{A}$ be an empty allocation
2 Find an independent set I of graph G (in any opportune way)
3 **foreach** $i \in I$ **do**
4 Agent i receives in $\mathcal{A}$ her most preferred object according to $\succ_i$ within O
5 Remove $\mathcal{A}(i)$ from O
6 Complete $\mathcal{A}$ (in any opportune way) and return $\mathcal{A}$

Proposition 1 *Algorithm 1 is* $\frac{|I|}{n}$*-approximate for* MAX-LEF*, where I is the independent set computed in step 2 of Algorithm 1.*

Proof. By construction, every member of I is LEF, and the largest number of LEF agents is $|N| = n$. □

Proposition 2 *The construction of I in Algorithm 1 (Step 2) can be done so that a polynomial time* $(\Delta(G)+1)^{-1}$*-approximation for* MAX-LEF *is produced, where $\Delta(G)$ is the maximum degree in G as introduced previously.*

Proof. The independent set is built as follows. I is initially empty and while $N \neq \emptyset$, do: choose $i \in N$, add i to I, and remove i and its neighbors from N. Since a node has at most $\Delta(G)$ neighbors, I is an independent set of size at least $n/(\Delta(G) + 1)$. Use Proposition 1 to get the expected ratio of $(\Delta(G) + 1)^{-1}$. □

The $(\Delta + 1)^{-1}$-approximation algorithm is long known for the MAXIMUM INDEPENDENT SET problem (that is, find an independent set of maximum cardinality) and slight improvements were proposed [43]. The following lemma shows that MAX-LEF shares exactly the same inapproximability results as MAXIMUM INDEPENDENT SET.

Lemma 2 *Any ρ-approximate algorithm for* MAX-LEF *is also a ρ-approximate algorithm for* MAXIMUM INDEPENDENT SET*.*

Proof. Suppose that we have a ρ-approximate algorithm for MAX-LEF and let us construct a ρ-approximate algorithm for MAXIMUM INDEPENDENT SET. Let $\mathcal{G}$ denote the graph of s vertices for which we look for an independent set of maximum size. Consider a set of s agents with identical preferences over a set of m objects. The agents are embedded in $\mathcal{G}$. Note that for any allocation $\mathcal{A}$ of objects to agents, the set of non-envious agents for $\mathcal{A}$ forms an independent set of $\mathcal{G}$. Indeed, a non-envious agent receives a better object in $\mathcal{A}$ than any of her neighbors. Therefore, each of her neighbors envies her, and no two neighbors can be non-envious. Furthermore, for each independent set I of $\mathcal{G}$, one can construct an allocation of objects to agents such that no agent of I is envious (one can use Algorithm 1). Therefore, the largest set of non-envious agents for any allocation of objects to agents is at least as large as the largest independent set in $\mathcal{G}$.

A ρ-approximate algorithm for MAX-LEF computes for this instance an allocation $\mathcal{A}$. Let I denote the set of non-envious agents for $\mathcal{A}$. As it was shown above, this set is an independent set. Furthermore, the size of I is at least ρ times the size of the largest set of non-envious agents for any allocation of objects to agents. Since the largest set of non-envious agents for any allocation of objects to agents is at least as large as the largest independent set in $\mathcal{G}$, I is a ρ-approximation for MAXIMUM INDEPENDENT SET in $\mathcal{G}$. □

MAXIMUM INDEPENDENT SET in general is **Poly-APX**-hard, meaning it is as hard as any problem that can be approximated to a polynomial factor. Lemma 2 implies that MAX-LEF is also **Poly-APX**-hard.

Proposition 3 MAX-LEF *is* **Poly-APX**-*hard.*

Thus, as long as we assume that $\mathbf{P} \neq \mathbf{NP}$, there is no constant approximation-ratio algorithm for MAX-LEF [5].

Interestingly, there are graph classes where the size of an independent set can be expressed as a fraction of n. Therefore, this fraction corresponds to the approximation ratio of Algorithm 1.

Proposition 4 *A polynomial time* 0.5-*approximate algorithm for* MAX-LEF *exists if the network is bipartite.*

Proof. Suppose that the network is a bipartite graph $(N_1, N_2; E)$. By definition both N_1 and N_2 are independent sets. If $|N_1| \geq |N_2|$, then run Algorithm 1 with $I := N_1$, otherwise run Algorithm 1 with $I := N_2$. Since $|I| = \max\{|N_1|, |N_2|\}$ and $|N| \leq 2\max\{|N_1|, |N_2|\}$, a polynomial time 0.5-approximation is reached. □

Proposition 4 can be easily extended to *k-partite graphs* (whose vertex set can be partitioned into k different independent sets), leading to a polynomial time k^{-1}-approximation algorithm.

Note that if the network admits a vertex cover C of size k, then Algorithm 1 with $I := N \setminus C$ provides a $(1 - k/n)$-approximate solution to MAX-LEF.

4.2 Optimizing degree of (non)-envy

Instead of simply counting the number of non-envious agents, we will now focus on a more subtle criterion, measuring the degree of envy among agents. This leads to the MAX-NE optimization problem (defined in Section 2) which consists in minimizing the *average degree of envy* $\mathcal{E}(\mathcal{A})$ (or equivalently maximizing the *average degree of non-envy* $\mathcal{NE}(\mathcal{A}) = 1 - \mathcal{E}(\mathcal{A})$). Before describing the algorithm, we first state the following lemma

Lemma 3 *Let $\mathcal{U}_n$ denote the uniform distribution over all matchings from n agents to n objects. Then we have $\mathbb{E}_{\mathcal{A}\sim\mathcal{U}_n}[\mathcal{NE}(\mathcal{A})] = \frac{5}{6} - o(1)$.*

Proof. In the following, by an abuse of notation we write $x, x' \sim O$ and $u, u' \sim [n]$ to mean that both x and x' (respectively both u and u') are drawn uniformly at random from O (respectively $[n]$).

$$\begin{aligned}
\mathbb{E}_{\mathcal{A}\sim\mathcal{U}_n}[\mathcal{E}(\mathcal{A})] &= \frac{1}{2|E|}\sum_{\{i,j\}\in E}\mathbb{E}_{\mathcal{A}\sim\mathcal{U}_n}[e(\mathcal{A},i,j)+e(\mathcal{A},j,i)] \\
&= \frac{1}{|E|(n-1)}\sum_{\{i,j\}\in E}\mathbb{E}_{x,x'\sim O}[\max(0, r_i(x) - r_i(x')) : x \neq x'] \\
&= \frac{1}{|E|(n-1)}\sum_{\{i,j\}\in E}\mathbb{E}_{u,u'\sim[n]}[\max(0, u-u') : u \neq u'] \\
&= \frac{1}{|E|(n-1)}|E| \cdot \mathbb{E}_{u,u'\sim[n]}[\max(0, u-u') : u \neq u'] \\
&= \frac{1}{(n-1)}\mathbb{E}_{u,u'\sim[n]}[\max(0, u-u') : u \neq u']
\end{aligned}$$

By the law of total expectation, we have:

$$\begin{aligned}
\mathbb{E}_{u,u'\sim[n]}[\max(0, u-u') : u \neq u'] &= \mathbb{E}_{u,u'\sim[n]}[u-u' : u > u'] \cdot P(u > u' : \\
&\quad u \neq u') + 0 \cdot P(u' > u : u \neq u') \\
&= \frac{1}{2}\mathbb{E}_{u,u'}[u-u' : u > u'] \\
&= \frac{1}{n(n-1)}\sum_{k=1}^{n-1}k(n-k) \\
&= \frac{n+1}{6}
\end{aligned}$$

Thus, $\mathbb{E}_{\mathcal{A}\sim\mathcal{U}_n}[\mathcal{NE}(\mathcal{A})] = 1 - \mathbb{E}_{\mathcal{A}\sim\mathcal{U}_n}[\mathcal{E}(\mathcal{A})] = 1 - \frac{1}{(n-1)} \cdot \frac{n+1}{6} = 1 - \frac{n+1}{6(n-1)} = \frac{5}{6} - o(1)$.

□

This tells us that with high probability, random matchings yield high degrees of non-envy. To get a deterministic algorithm based on this idea, we apply a standard derandomization technique. In our algorithm (Algorithm 2), at each step i, agent i receives one of the remaining unallocated objects. $\mathcal{A}_i$ denotes the partial allocation built up to step i where each agent $j \leq i$ is assigned an object. This object is chosen so as to minimize the conditional expectation of $\mathcal{E}$ (line 5). $\mathcal{A}_i^x$ is one of the possible allocations built from $\mathcal{A}_{i-1}$ by allocating x to agent i. We will show below that this conditional expectation can be computed efficiently.

Algorithm 2:

```
A_0 is an empty allocation
for each agent i ∈ N do
    U is the set of unassigned objects in A_{i−1}
    for each object x ∈ U do
        A_i^x ← A_{i−1} ∪ {(i, x)}
        v_x ← E_{A∼U_n} [E(A) : A_i^x ⊆ A]
    x* ← arg min_{x∈U} v_x
    A_i ← A_{i−1} ∪ {(i, x*)}
```

Proposition 5 *Algorithm 2 is a polynomial-time* $\frac{5}{6} - o(1)$ *approximation algorithm for* MAX-NE*.*

Proof. First, by standard arguments of the derandomization method (similar to e.g. page 132 of Vazirani's book on approximation algorithm [48]) together with Lemma 3, we will show that this algorithm outputs an allocation $\mathcal{A}_N$ such that $\mathcal{NE}(\mathcal{A}_N) \geq \frac{5}{6} - o(1)$. By design we have $\mathcal{NE}(\mathcal{A}^*) \leq 1$ where $\mathcal{A}^*$ is the assignment which maximizes the degree of non-envy, so the approximation ratio holds. Let us show that $\mathcal{E}(\mathcal{A}_N) \leq \frac{1}{6} + o(1)$. Because $\mathcal{A}_N$ is not a partial allocation, $\mathcal{E}(\mathcal{A}_N) = \mathbb{E}_{\mathcal{A}\sim\mathcal{U}_n}\left[\mathcal{E}\left(\mathcal{A}\right) : \mathcal{A}_N \subseteq \mathcal{A}\right]$, so showing by induction that for all $i \in N$, $\mathbb{E}_{\mathcal{A}\sim\mathcal{U}_n}\left[\mathcal{E}\left(\mathcal{A}\right) : \mathcal{A}_i \subseteq \mathcal{A}\right] \leq \frac{1}{6} + o(1)$ will conclude this part of the proof. At iteration $i = 1$, we have $\mathbb{E}_{\mathcal{A}\sim\mathcal{U}_n}\left[\mathcal{E}\left(\mathcal{A}\right) : \mathcal{A}_{i-1} \subseteq \mathcal{A}\right] = \mathbb{E}_{\mathcal{A}\sim\mathcal{U}_n}\left[\mathcal{E}\left(\mathcal{A}\right)\right] = \frac{1}{6} + o(1)$ by Lemma 3. At iteration $i > 1$ of the algorithm, by the law of total expectation we have

$$\begin{aligned}\mathbb{E}_{\mathcal{A}\sim\mathcal{U}_n}\left[\mathcal{E}\left(\mathcal{A}\right) : \mathcal{A}_{i-1} \subseteq \mathcal{A}\right] &\geq \min_x \mathbb{E}_{\mathcal{A}\sim\mathcal{U}_n}\left[\mathcal{E}\left(\mathcal{A}\right) : \mathcal{A}_{i-1} \cup \{(i,x)\} \subseteq \mathcal{A}\right] \\ &= \mathbb{E}_{\mathcal{A}\sim\mathcal{U}_n}\left[\mathcal{E}\left(\mathcal{A}\right) : \mathcal{A}_i \subseteq \mathcal{A}\right]\end{aligned}$$

Thus, the conditional expectations are non increasing. Therefore, at iteration $i = N$ we have $\mathcal{E}(\mathcal{A}_N) \leq \frac{1}{6} + o(1)$.

Next, to show that the algorithm runs in polynomial time, we need to bound the computation time of v_x. If $\mathcal{A}$ is a partial allocation, define $P(\mathcal{A}, l)$ as the set of goods that agent l can own without violating $\mathcal{A}$. For example, if $\mathcal{A}$ is a complete allocation, $P(\mathcal{A}, l) = \mathcal{A}(l)$ and if $\mathcal{A} = \{\}$, then $P(\mathcal{A}, l) = O$. First note that due to the fact that the expectation operator is linear, v_x can be calculated as a sum of conditional expectations $\frac{1}{2|E|}\sum_{\{l,h\}\in E}\mathbb{E}_{\mathcal{A}\sim\mathcal{U}_n}\left[e(\mathcal{A}, l, h) + e(\mathcal{A}, h, l) : \mathcal{A}_i^x \subseteq \mathcal{A}\right]$. Next, note that for any $l, h \in N$ the expectation $\mathbb{E}_{\mathcal{A}\sim\mathcal{U}_n}[e(\mathcal{A}, l, h) : \mathcal{A}_i^x \subseteq \mathcal{A}]$is equal to $\frac{1}{|Z_{l,h}|\cdot(n-1)}\sum_{(a,b)\in Z_{l,h}} \max(0, r_l(a) - r_l(b))$ where $Z_{l,h} = \left\{(a, b) \in P(l, \mathcal{A}_i^x) \times P(h, \mathcal{A}_i^x) : a \neq b\right\}$. The computation of v_x can thus be done in $O(n^4)$. $\square$

5 Location and allocation

This section is dedicated to DEC-LOCATION-LEF. The problem asks whether there exists an assignment of agents to positions of the graph as well as an assignment of objects to agents such that the resulting allocation is locally envy-free. The following theorem shows that this problem is computationally challenging.

Theorem 7 DEC-LOCATION-LEF *is* **NP**-*complete.*

Proof. The reduction is from problem INDEPENDENT SET which is **NP**-complete [32] and can be defined as follows. An instance $\mathcal{I}$ of INDEPENDENT SET is described by an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and a positive integer $k \leq |\mathcal{V}|$, the question is whether there exists an independent set $I \subseteq \mathcal{V}$ of size k. Let s denote the size of $\mathcal{V}$, and $\mathcal{V} = \{v_1, \cdots, v_s\}$.

We construct an instance $\mathcal{J}$ of DEC-LOCATION-LEF as follows. The set of objects is $O = Q \cup T$, where $Q = \{q_1, \cdots, q_{s-k}\}$ and $T = \{t_1, \cdots, t_k\}$. The set of agents is $N = \{X_1, \ldots, X_{s-k}\} \cup \{L_1, \ldots, L_k\}$. Let Q_{-i} denote the set $Q \setminus \{q_i\}$, and let $Q^{\uparrow}_{-i}$, $Q^{\uparrow}$ and $T^{\uparrow}$ denote partial orders over Q_{-i}, Q and T, respectively, where objects are ranked by increasing order of indices. Preferences are as follows:

- $X_i : q_i \succ Q^{\uparrow}_{-i} \succ T^{\uparrow}$
- $L_j : T^{\uparrow} \succ Q^{\uparrow}$

Finally, the network is $G = \mathcal{G} = (\mathcal{V}, \mathcal{E})$.

We claim that we can place agents in $\mathcal{G}$ and allocate them objects such that there is no envy in $\mathcal{J}$ if, and only if, $\mathcal{G}$ contains an independent set of size k in $\mathcal{I}$.

Assume that I is an independent set of size k in $\mathcal{G}$. We can assume without loss of generality that $I = \{v_1, \ldots, v_k\}$. We construct $\mathcal{A}$ and $\mathcal{L}$ as follows. If $v_i \in I$ then $\mathcal{L}(L_i) = v_i$ and $(L_i, t_i) \in \mathcal{A}$. Otherwise, agents are placed arbitrarily on $\mathcal{G}$ and receive their most preferred object (i.e., $(X_i, q_i) \in \mathcal{A}$). Thus, no agent X_i will envy one of her neighbors. Furthermore, no two vertices $\mathcal{L}(L_i)$ and $\mathcal{L}(L_j)$ are neighbors in $\mathcal{G}$. Therefore, no agent L_j will envy one of her neighbors, say X_i, who receives object q_i that is less preferred by agent L_j to object t_i.

Assume now that there exists $\mathcal{L}$ and $\mathcal{A}$ such that no agent envies one of her neighbors in $\mathcal{G}$ when her position is defined by $\mathcal{L}$ and her assignment is defined by $\mathcal{A}$. For any L_i and L_j, either $\mathcal{A}(L_i) \succ_{L_j} \mathcal{A}(L_j)$ or $\mathcal{A}(L_j) \succ_{L_i} \mathcal{A}(L_i)$ holds since L_i and L_j have the same preferences. Therefore, $\mathcal{L}(L_i)$ and $\mathcal{L}(L_j)$ cannot be neighbors in $\mathcal{G}$ since otherwise either L_i would envy L_j or L_j would envy L_i. Hence, $\{\mathcal{L}(L_1), \ldots, \mathcal{L}(L_k)\}$ forms an independent set of size k in $\mathcal{G}$. □

Interestingly, the above reduction also holds when $\mathcal{A}$ is fixed, i.e. the allocation of objects to agents is imposed by the problem.

We shall extend the polynomial time result obtained for DEC-LEF on networks of degree at least $n - 2$.

We claim that two agents having the same top object must be neighbors in $\overline{G}$ in any YES-instance. Indeed, otherwise if one of them obtains her most preferred object then the other will be envious, and if this object is assigned to one of her common neighbor in $\overline{G}$ (see Observation 2) then this vertex will have a degree at most $n-3$ in G. Therefore, one can focus on $\mathbb{L}_{\succ}$, defined as the set of location functions such that each pair of agents having the same top object are neighbors in $\overline{G}$ (or equivalently, not neighbors in G).

If an instance contains three (or more) agents with the same top object then it must be a NO-instance since each vertex in $\overline{G}$ has degree at most 1 and no three agents can be neighbors in $\overline{G}$. The following lemma shows that the location functions of $\mathbb{L}_{\succ}$ are all equivalent for the search of an LEF allocation.

Lemma 4 *If $\mathcal{A}$ is an LEF allocation for some $\mathcal{L}$, and $\mathcal{A}$ is Pareto-optimal among the LEF allocations determined with respect to $\mathcal{L}$, then $\mathcal{A}$ is also LEF for any location function of $\mathbb{L}_{\succ}$.*

Proof. First of all, $\mathcal{L}$ must belong to $\mathbb{L}_{\succ}$ for $\mathcal{A}$ to be LEF. Let $\mathcal{L}'$ be another function of $\mathbb{L}_{\succ}$. Since any pair of agents having the same top object should be neighbors in G for any location function of $\mathbb{L}_{\succ}$, they have the same set of neighbors in G in both $\mathcal{L}$ and $\mathcal{L}'$. Therefore, if none of these agents envies one of her neighbors under $\mathcal{L}$ with allocation $\mathcal{A}$, then they neither envy their neighbors under $\mathcal{L}'$ with allocation $\mathcal{A}$.

Let i be an agent who is the only one to rank some object o at the first position in her preferences. On one hand, if $\mathcal{L}(i)$ is a vertex of degree $n-1$ then Observation 2 implies that she must receive o. On the other hand, if $\mathcal{L}(i)$ is a vertex of degree $n-2$ and j is the unique neighbor of i in $\overline{G}$ then Observation 2 implies that o is assigned either to i or to j. But j must also be the unique agent to have some object o' ranked first in her preferences, where $o \neq o'$, because otherwise the other agent who ranks o' first in her preference would not be the neighbor of j under $\mathcal{L}(i)$. Therefore, either agent i or j must receive o'. Since by hypothesis $\mathcal{A}$ is Pareto-optimal among the LEF allocations determined with respect to $\mathcal{L}$, o must be assigned to i and o' must be assigned to j. All in all, agent i must receive her top object in $\mathcal{A}$ and envies none of the other agents. Therefore, no agent will envy her neighbor under $\mathcal{L}'$ with allocation $\mathcal{A}$. □

In order to solve DEC-LOCATION-LEF, one can compute a function $\mathcal{L}$ of $\mathbb{L}_{\succ}$ by assigning the agents having the same top object to vertices connected in $\overline{G}$, and by assigning the other agents arbitrarily. If such a location does not exist then the instance is a NO-instance. Once $\mathcal{L}$ is fixed, one can use the algorithm presented in Theorem 2 to compute an LEF allocation if such an allocation exists. If an LEF allocation $\mathcal{A}$ is returned then the algorithm returns $\mathcal{L}$ and $\mathcal{A}$. Otherwise, we know by Lemma 4 that no function in $\mathbb{L}_{\succ}$ can lead to an LEF allocation (if an LEF allocation for $\mathcal{L}$ had existed, then an LEF allocation which is Pareto-optimal among the LEF allocations determined with respect to $\mathcal{L}$ would necessarily exist), and the algorithm returns *false*. This algorithm clearly runs in polynomial time, leading to the following theorem:

Theorem 8 DEC-LOCATION-LEF *in graphs of minimum degree $n-2$ is solvable in polynomial time.*

6 The likelihood of locally envy-free allocations

In order to better understand the impact of the structure of the graph on local envy-freeness, we run some experiments where we investigate the influence of different characteristics of the network. In particular, we observe the impact of the degree of the nodes in all the problems that we have studied and the behavior of specific classes of graphs close to real networks in the existence of locally envy-free allocations.

6.1 Impact of the degree of the nodes

In this subsection we generate random instances of our decision and optimization problems and use mixed integer linear program formulations to compute the optimal solutions of these instances. We build on the ones proposed by Dickerson et al. [25] (which address envy-freeness and the minimization of maximum pairwise envy among any two agents [41], in a context of additive utilities with several goods per agent). To fit our setting, we adapt it so as to account for graph constraints, the constraint that exactly one object per agent is to be allocated, and strict ordinal (linear) preferences over these objects. We further design three variants, two where the objective functions correspond to MAX-LEF and MAX-NE, and another one where the locations of agents on the graph are treated as decision variables, to address the more challenging DEC-LOCATION-LEF. The interested reader may find this MIP formulation in Appendix A.

For these experiments, we generate random regular graphs of degree k with 8 vertices for k ranging from 1 to 7. We rely on the graph generator of the Python module NetworkX [35], which produces random regular graphs using the algorithm of Steger and Wormald [45]. Agents' preferences are randomly drawn from impartial culture [15, 33], that is, all possible preference orders are equally likely and chosen independently. Table 3 shows the results (averaged over 1,000 runs). The entry LEF gives the likelihood of picking a *yes*-instance of the LEF problem, whereas Loc-LEF gives this likelihood for DEC-LOCATION-LEF. On line MAX-LEF, we report the number of remaining envious agents after solving the MAX-LEF optimization problem, and MAX-NE gives the average degree of non-envy after solving MAX-NE optimization problem. Finally, the entry MMPE corresponds to the "classical" minimization of maximum pairwise envy (MMPE) of Lipton et al. [41] which in our context can be interpreted as the maximum number of agents envied by any agent.

Degree	1	2	3	4	5	6	7
LEF	1	0.72	0.22	0.05	0.02	<0.01	<0.01
Loc-LEF	1	1	1	0.92	0.49	0.07	<0.01
MAX-LEF	0	0.28	0.93	1.52	1.95	2.44	2.78
MAX-NE	1	0.99	0.99	0.99	0.98	0.98	0.98
MMPE	0	0.28	0.83	1.19	1.42	1.69	1.91

Table 3: Experimental results for our decision and optimization problems in 1,000 instances with 8 agents and graphs of regular degree. The LEF and Loc-LEF lines refer to the likelihood of existence of an LEF allocation in DEC-LEF and DEC-LOCATION-LEF problems. The MAX-LEF and MAX-NE lines refer to the number of LEF agents in average and to the average degree of non-envy, respectively, after optimization. The MMPE line gives the maximum number of agents envied by any agent, after optimization.

Discussion about the results. The first question that we address is how the likelihood to pick a positive instance of DEC-LEF evolves, under impartial culture. It must clearly decrease: in the extreme case of a complete graph, recall that all agents must have a different preferred item, which occurs with a low probability, i.e. with probability equal to $n!/n^n$.

In fact, asymptotically (as the number of agents grows to infinity), it can be shown that the likelihood to pick a positive instance is negligible, as soon as the degree of the graph is above a fraction of $1/e$ ($\simeq 0.36$) of the overall number of nodes. The following proposition formally states this:

Proposition 6 *Assume agent preferences are drawn from impartial culture. Suppose that $\delta(G) \geq cn$, where $c > e^{-1}$ is some constant. Then almost surely G is not locally envy-free.*

Proof. Consider a fixed allocation $\mathcal{A}$. For any agent i, define $NE(i)$ as the event that i does not envy her neighbors. Clearly, since preference lists are completely independent, the events $NE(i)$ and $NE(j)$ are independent for all $i \neq j$. We note also that as preference lists are random permutations, $P\left(NE(i)\right) = \frac{1}{deg_G(i)+1}$ for every agent i. Thus, clearly

$$P(\mathcal{A} \text{ is locally envy-free}) \leq \left(\frac{1}{\delta(G)+1}\right)^n$$

Let X be the number of locally envy-free allocations. It follows that $\mathbb{E}[X] \leq n!\left(\frac{1}{cn}\right)^n \leq \sqrt{2\pi n}\left(\frac{n}{e}\right)^n e^{\frac{1}{12n}}\left(\frac{1}{cn}\right)^n$, by Stirling's formula. Thus, $\mathbb{E}[X] \leq 3\sqrt{n}\left(\frac{1}{ec}\right)^n$. Thus, by Markov's inequality, the probability that there is an LEF allocation is at most $3\sqrt{n}\left(\frac{1}{ec}\right)^n = o(1)$, since $ce > 1$. □

On the other hand, for graphs of small degrees, it is often the case that an LEF can be found. The question is thus how this drop will occur. Our experiments, displayed in Table 3, suggest that this decrease is sharp.

Now let us turn our attention to DEC-LOCATION-LEF. The ability to allocate agents on the network gives the central authority some extra-power when it comes to finding an LEF. However, note that this power heavily depends on the structure of the graph (for instance, it is useless when the graph is complete, as all the different ways to label the graph with agents are isomorphic). The entry Loc-LEF of Table 3 shows that this power can be significant: the likelihood to pick an instance where an LEF allocation exists in that context remains above 90% until degree 4, while it was as low as 5.5% in the basic problem.

Although an LEF allocation may not exist, positive results are obtained regarding the measures we optimize. Even with a complete graph, it is on average possible to allocate items so as to make envious only about a third of the agents, and such that no agent envies more than two other agents in our instance with 8 agents. Indeed, when the degree of the graph is 7 (i.e. the graph is complete), the average number of envious agents (MAX-LEF) is 2.78 and the average maximum number of agents envied by any agent (MMPE) is 1.91.

6.2 Empirical existence of LEF allocations in realistic networks

In this subsection, we are especially concerned with the frequency of positive instances of DEC-LEF, that is how often a locally envy-free allocation exists, and how many LEF allocations there are when they exist. We conduct experiments in more realistic settings, in particular by considering domain restrictions for preferences, and graph structures which may be induced by those agents' individual preferences.

Like in the first subsection, we run 1,000 instances with 8 agents. The linear preferences of the agents are generated either from impartial culture (IC), with no restriction of domain, or following two different distributions for preferences restricted to the *single-peaked* domain [14]. Let us recall that a preference order $\succ_i$ is single-peaked with respect to an axis $>^O$ over the objects if there exists a unique peak object $x^* \in O$ such that for every couple of objects a and b, $x^* >^O a >^O b$ implies that $x^* \succ_i a \succ_i b$ and $a >^O b >^O x^*$ implies that $x^* \succ_i b \succ_i a$. In our experiments, the single-peaked preferences are generated from the *single-peaked uniform* culture (SP-U), i.e., they are uniformly drawn from the urn containing all single-peaked rankings with respect to a given axis over the objects, or from the *single-peaked uniform peak* culture (SP-UP), i.e., they are generated by uniformly drawing a peak alternative on a given axis over the objects and then by iteratively choosing the next preferred alternatives with equal probability either on the left of the peak in the axis, or on the right.

We restrict ourselves to specific classes of graphs, and more precisely to Barabási-Albert random graphs [10], graphs with homophily, and graphs complementary to graphs with homophily (we refer to them as graphs with *het-*

erophily). These graphs are supposed to be closer to real networks than simple random graphs.

The Barabási-Albert graphs are typical *scale-free* networks. The scale-free property is usually found in real networks and has been formulated in [10]. A network is scale-free if the degree of its vertices follows a power-law distribution, that is the fraction of vertices of degree k is proportional to $k^{-\gamma}$, for some constant γ. The main observable feature on a scale-free network is that it contains many *hubs*, that are nodes with high degree. The Barabási-Albert random graphs are scale-free because the degree of their nodes follows a power-law distribution with degree exponent γ equal to three. In such random graphs, the network is iteratively constructed by adding to a subgraph a new node which is connected with higher probability to high degree nodes, following a preferential attachment mechanism. More precisely, given a subgraph G' defined on a subset of vertices $N' \subseteq N$, a new graph G'' is constructed by adding a new node $i \in N \setminus N'$ and a new edge connecting i and any node $j \in N'$ with probability $p_j = \frac{deg_{G'}(j)}{\sum_{i' \in N'} deg_{G'}(i')}$.

A network respects *homophily* if two "similar" nodes tend to be connected in the graph. In the context where agents are embedded in a network, two agents can be considered similar if they have close preferences over the set of objects. We generate graphs with homophily by following a protocol adapted for taking into account ordinal preferences: the more the agents agree on pairwise comparisons of the objects, the more likely they are connected. More precisely, two agents i and j are connected via edge $\{i, j\}$ in G with probability equal to $q_{ij} = |\{(a, b) \in O^2 : a \succ_i b \text{ and } a \succ_j b\}|/\binom{n}{2}$. Intuitively, in this model, the probability of connection between two agents is inversely proportional to the Kendall-Tau distance between their respective preference rankings. In particular, two agents with exactly the same preferences are necessarily connected.

The results concerning the frequency of existence of an LEF allocation are presented in Figure 7.

Observe that the likelihood of finding an LEF allocation in a graph with homophily is extremely low. Indeed, the closer the preferences of two agents, the more likely they are to be connected in the network with homophily. Therefore, it appears natural that finding an LEF allocation in such instances is difficult. On the contrary, when this is the complementary graph, i.e., the non-envy graph $\overline{G}$, that respects homophily, the likelihood of finding an LEF allocation is clearly higher. Naturally, there are more LEF allocations when the complementary graph respects homophily because two agents with very different preferences are more likely to be connected than two agents with similar preferences. Therefore, it should be easier to find an LEF allocation in such graphs. The likelihood of existence of an LEF allocation in such graphs is even higher for single-peaked profiles: an LEF allocation exists in 30% of the instances under impartial culture whereas this frequency is around 40% for SP-U profiles and more than 60% for SP-UP profiles. The significant increase in the frequency of existence of LEF allocations for single-peaked profiles where the

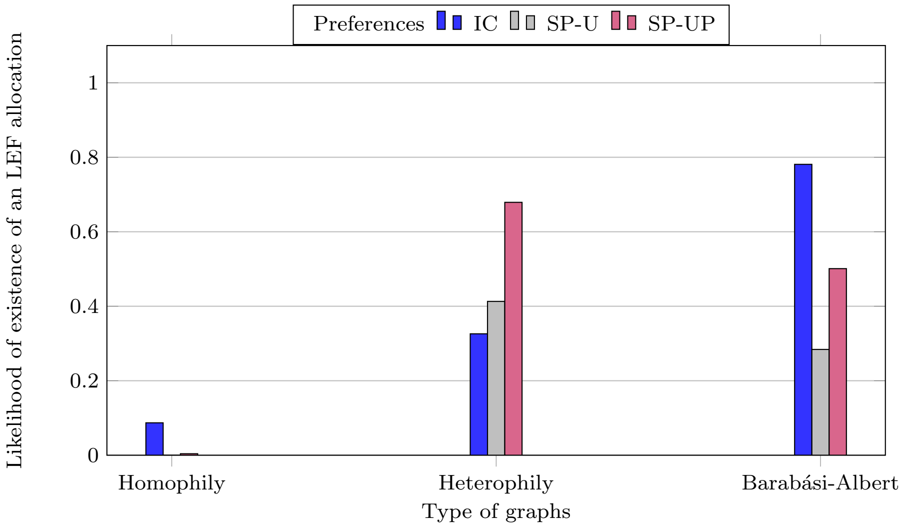

Fig. 7: Likelihood of existence of an LEF allocation in DEC-LEF problem for different classes of graphs (with homophily, with heterophily or of type Barabási-Albert) and different types of ordinal preferences (IC, SP-U or SP-UP) in 1,000 generated instances with 8 agents

top object of each agent is uniformly drawn (SP-UP) is due to the fact that preference orders that are single-peaked with the same top object (or peak) are very close regarding the number of common pairwise comparisons of objects. Therefore, in single-peaked profiles, two agents having the same top object are very likely to be connected in graphs respecting homophily, and thus are very likely to not be connected in graphs with heterophily. On the contrary, by construction of a single-peaked preference order, two preference orders with different peak objects, and especially two peak objects that are far from each other on the single-peaked axis over the objects, tend to agree on a very few number of pairwise comparisons of objects. This is typically the case for the two extreme points of the single-peaked axis which induce unique preference orders that are completely reversed; two agents having these respective preferences are necessarily connected in graphs with heterophily. Preference orders with different peak objects appear more frequently in single-peaked profiles SP-UP than in single-peaked profiles SP-U. Indeed, in the SP-UP distribution, the probability to pick a preference order with a given peak object $o \in O$ is equal to $1/n$ for every object o whereas, for instance, the probability of picking the unique single-peaked preference order with an "extreme" object (in the single-peaked axis) as its peak is equal to $1/2^{n-1}$ in the SP-U distribution.

Concerning the Barabási-Albert graphs, the likelihood of finding an LEF allocation is high under impartial culture (around 80%), low for SP-U profiles (around 30%) and medium for SP-UP profiles (around 50%). For these graphs, contrary to graphs where the complement respects homophily, the likelihood of finding a locally envy-free allocation is significantly higher in profiles generated uniformly (IC). This can be explained by the fact that the preferences of the agents are less correlated and thus globally the agents do not desire the same objects. Similarly, as with graphs whose complement respects homophily, the likelihood of finding an LEF allocation is higher in SP-UP profiles than in SP-U profiles because the preferences of the agents are more diverse.

Regarding the number of LEF allocations, the results are presented in Table 4. The numbers are given in average without counting the negative instances for existence in order to have a clearer idea and not being noised by the numerous instances with no LEF allocations. The instances are the same as those considered for testing the likelihood of existence. Recall that the total number of possible house allocations for instances with 8 agents is equal to $8! = 40,320$.

	IC	SP-U	SP-UP
Homophily	4.26	0	1.5
Heterophily	9.86	267.08	146.40
Barabasi-Albert	7131.24	18006.51	11744.64

Table 4: Number of LEF allocations in positive instances of DEC-LEF for different classes of graphs with 8 agents

In accordance with the likelihood of existence which is very weak for networks with homophily, the number of locally envy-free allocations is also extremely low in such graphs. For the other types of graphs, the number of LEF allocations is rather high, especially for Barabási-Albert graphs. This is due to the very low density of Barabási-Albert graphs compared to the other graphs. Globally, the number of LEF allocations is smaller under impartial culture, even for Barabási-Albert graphs for which the frequency of LEF existence is the highest under impartial culture. Moreover, the number of LEF allocations is the highest for SP-U profiles. This may be surprising because the likelihood of existence of an LEF allocation is the lowest for SP-U profiles (see Figure 7). This phenomenon may be explained by the fact that the few instances with existence of LEF allocations in SP-U profiles may exhibit opportune configurations where the LEF allocations are very numerous.

7 Conclusion and future work

We have studied different aspects of local envy-freeness in house allocation settings. First of all, deciding whether a locally envy-free allocation exists in a given instance is computationally hard even for very simple and sparse graphs. Nevertheless, we were able to provide some tractable cases according to the network topology. See Table 1 for the details of the complexity results and polynomial cases, with respect to some parameters of the graph. Interestingly, DEC-LEF is solvable in polynomial-time in graphs of degree at least $n-2$. This case is very interesting because it relies on meaningful envy-graphs. Indeed, the graphs with degree at least $n-2$ refer to the case where the non-envy graph is of degree at most 1, and thus includes the case where the non-envy graph is composed of couples of agents within which there is no reason for envy to exist.

We have also investigated an optimization perspective, and have tried to maximize the number of LEF agents or minimize the average envy with respect to specific degrees of envy. We have provided for both cases approximation algorithms. Due to its connection with MAXIMUM INDEPENDENT SET (cf. Lemma 2), significant improvements for MAX-LEF are unlikely.

In a third direction, we have studied the power of the central authority by assuming that, given the structure of the network, she can assign in addition to the items to the agents, the agents to the locations on the graph. This problem can be understood as assigning tasks to workers as well as time slots (see Example 1). Although hard in general, this problem is solvable in polynomial time for the interesting case of graphs of degree at least $n-2$.

Finally, the experiments confirm the intuition that the likelihood of finding a locally envy-free allocation is higher in sparse graphs. But they also highlight the fact that for some graphs close to real social networks, such as non-envy graphs with neighbors having similar preferences or scale-free networks, the probability of existence of an LEF allocation is rather high as well as the number of such allocations.

There are several interesting future directions to explore. We give below some preliminary thoughts on the ones we find the most stimulating.

Constraints on the allocation. Two other relevant challenges related to DEC-LEF are: Given a *partial* allocation of the objects, can a full LEF allocation be found? Given some forbidden object-agent pairs, can an LEF allocation be found?

Pareto-efficiency and LEF. This paper leaves aside efficiency concerns, except for the fact that objects should not be wasted. A natural question is how LEF requirements interplay with Pareto-efficiency. As we have seen already, as soon as the network is not complete, an LEF allocation is not necessarily Pareto-efficient. More interestingly, it is also not the case that at least one of the LEF allocation is Pareto. This can be seen on the example depicted in Figure 8.

This instance admits a single LEF allocation $\{(1,a),(2,b),(3,d),(4,c)\}$, which is Pareto-dominated by $\{(1,c),(2,b),(3,d),(4,a)\}$, which is *not* LEF (agent 3 would envy agent 4). Based on these remarks, it would be interesting to study Pareto-efficient allocations within the set of LEF allocations (a notion which also emerged in Lemma 4, but remains to be studied in depth).

Oriented graphs. A natural extension of DEC-LEF is to consider a network modeled with a directed graph. An arc (u,v) indicates that u possibly envies v, but it does not indicate that v possibly envies u, unless the arc (v,u) is also

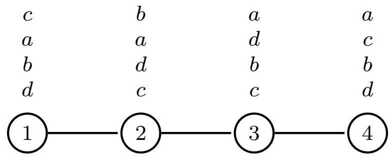

Fig. 8: The only LEF allocation is not Pareto

present. In this directed case, an allocation $\mathcal{A}$ is said to be LEF if $\mathcal{A}(j) \not\succ_i \mathcal{A}(i)$ for every arc (i, j). It is not difficult to see that DEC-LEF is **NP**-complete in this directed case (use the proof of Theorem 1 where each edge $\{u, v\}$ is replaced by the arcs (u, v) and (v, u)). Interestingly, the directed variant of DEC-LEF can be solved efficiently in *directed acyclic graphs* (DAGs). Indeed, if the network is a DAG, then an LEF allocation must exist, and it can be computed in polynomial time. In fact, a DAG has at least one *source*, i.e. a vertex with indegree 0. If a source of a DAG is deleted, then we get a (possibly empty) DAG. The algorithm computing an LEF allocation works as follows: while the network is non empty, find a source s, allocate s her most preferred object $o_s \in O$, remove o_s from O, and delete s. The algorithm also guarantees a Pareto-optimal allocation and mimics a serial dictatorship [47].

DAGs have also been considered by Bredereck et al. [19, Observation 3] but their algorithm is different because giving nothing to an agent is not allowed in our setting.

Note that DAGs actually characterize exactly those graphs guaranteeing LEF to exist (if a cycle exists, simply set the preferences of all agents to be exactly the same within the cycle to get a no-instance). But this leaves other interesting questions open: for instance, are there other natural classes of graphs admitting polynomial time algorithms for DEC-LEF in oriented graphs?

Domain restrictions. There is a long tradition in social choice to consider domain restrictions on agents' preferences to obtain positive results. This would be natural to study our setting under such assumptions. For example, we can fix the number of different rankings. To take a concrete question, how difficult DEC-LEF and DEC-LOCATION-LEF are when there are only two categories of agents: those with ranking $\succ_1$ on the objects and those with ranking $\succ_2$? More generally, can well-known domain restrictions, such as single-peakedness, be useful? Since the relevance of this domain restriction in the context of house allocation has recently been emphasized [9, 24], this might be an interesting road to pursue. As a first step in that direction, we have conducted experiments which provide some insights regarding the likelihood of locally envy-free allocations in this domain.

A Appendix: MIP formulation for dec-location-LEF

We describe the MIP formulation used to solve the DEC-LOCATION-LEF problem. We are given a set of objects O, a set of agents N equipped with preferences over those objects (for the ease of exposure we refer here to $r_{i,o}$ as the rank of object o in the preference order of agent i), and a graph $G = (V, E)$.

Together with the real valued decision variable e, which will be used to express the envy bound we try to minimize, we make use of the following (binary) decision variables:

- $x_{i,o}$: agent i holds object o
- $l_{i,p}$: agent i is located on node p
- $s_{i,j,o}$: agent i sees that agent j holds object o

We first express that each agent must receive exactly one object, and that each object must be assigned to exactly one agent (constraints (1) and (2)). Similarly, each agent must be assigned to a single node of the network, and each node must have a single agent assigned (constraints (3) and (4)).

$$\forall i \in N : \textstyle\sum_{o \in O} x_{i,j} = 1 \tag{1}$$
$$\forall o \in O : \textstyle\sum_{i \in N} x_{i,o} = 1 \tag{2}$$
$$\forall i \in N : \textstyle\sum_{p \in V} l_{i,p} = 1 \tag{3}$$
$$\forall p \in V : \textstyle\sum_{i \in N} l_{i,p} = 1 \tag{4}$$

When agent i is located on a node p connected to a node q where agent j holds o, i sees that j holds o:

$$\forall i, j \in N, \forall \{p, q\} \in E : l_{i,p} + l_{j,q} + x_{j,o} - 2 \leq s_{i,j,o} \tag{5}$$

Finally, we try to minimize the amount of envy between any pair of agents (MMPE), which is expressed by setting, together with the objective function $\min e$, constraint (6):

$$\forall i, j \in N : \sum_{o \in O} r_{i,o} \times s_{i,j,o} - \sum_{o \in O} r_{i,o} \times x_{i,o} \leq e \tag{6}$$

Note that in the case of DEC-LOCATION-LEF, we are only interested in whether we can find a solution which sets the envy bound e at 0, i.e., whether an LEF allocation exists.

Acknowledgements This work is partially supported by the ANR project 14- CE24-0007-01- CoCoRICo-CoDec. We thank the reviewers of the conference and journal versions for their useful comments.

References

1. Abdulkadiroğlu, A., Sönmez, T.: Random serial dictatorship and the core from random endowments in house allocation problems. Econometrica **66**(3), 689–701 (1998)
2. Abebe, R., Kleinberg, J., Parkes, D.C.: Fair division via social comparison. In: Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-17), pp. 281–289. ACM, São Paulo, Brazil (2017)
3. Amanatidis, G., Birmpas, G., Markakis, V.: Comparing approximate relaxations of envy-freeness. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI-18), pp. 42–48. International Joint Conferences on Artificial Intelligence, Stockholm, Sweden (2018)
4. Aspvall, B., Plass, M.F., Tarjan, R.E.: A linear-time algorithm for testing the truth of certain quantified boolean formulas. Information Processing Letters **8**(3), 121–123 (1979)

5. Ausiello, G., Crescenzi, P., Gambosi, G., Kann, V., Marchetti-Spaccamela, A., Protasi, M.: Complexity and approximation: Combinatorial optimization problems and their approximability properties. Springer Science & Business Media (2012)
6. Aziz, H., Bouveret, S., Caragiannis, I., Giagkousi, I., Lang, J.: Knowledge, fairness, and social constraints. In: Proceedings of the 32nd Conference on Artificial Intelligence (AAAI-18), pp. 4638–4645. AAAI Press, New Orleans, Louisiana, USA (2018)
7. Aziz, H., Hougaard, J.L., Moreno-Ternero, J.D., Østerdal, L.P.: Computational aspects of assigning agents to a line. Mathematical Social Sciences **90**, 93–99 (2017)
8. Aziz, H., Schlotter, I., Walsh, T.: Control of fair division. In: Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-16), pp. 67–73. IJCAI/AAAI Press, New York, NY, USA (2016)
9. Bade, S.: Matching with single-peaked preferences. Journal of Economic Theory **180**, 81–99 (2019)
10. Barabási, A.L., Albert, R.: Emergence of scaling in random networks. Science **286**(5439), 509–512 (1999)
11. Bei, X., Qiao, Y., Zhang, S.: Networked Fairness in Cake Cutting. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-17), pp. 3632–3638. International Joint Conferences on Artificial Intelligence, Melbourne, Australia (2017)
12. Berman, P., Karpinski, M., Scott, A.D.: Approximation hardness of short symmetric instances of MAX-3SAT. Electronic Colloquium on Computational Complexity (ECCC) (49) (2003)
13. Bilò, V., Caragiannis, I., Flammini, M., Igarashi, A., Monaco, G., Peters, D., Vinci, C., Zwicker, W.S.: Almost envy-free allocations with connected bundles. In: Proceedings of the 10th Conference on Innovations in Theoretical Computer Science (ITCS-19). San Diego, California, USA (2019)
14. Black, D.: On the rationale of group decision-making. Journal of Political Economy **56**(1), 23–34 (1948)
15. Black, D., Newing, R.A., McLean, I., McMillan, A., Monroe, B.L.: The theory of committees and elections. Springer (1958)
16. 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-17), pp. 135–141. International Joint Conferences on Artificial Intelligence, Melbourne, Australia (2017)
17. Bouveret, S., Chevaleyre, Y., Maudet, N.: Fair allocation of indivisible goods. In: F. Brandt, V. Conitzer, U. Endriss, J. Lang, A.D. Procaccia (eds.) Handbook of Computational Social Choice, chap. 12, pp. 284–310. Cambridge University Press (2016)
18. Bouveret, S., Lemaître, M.: Characterizing conflicts in fair division of indivisible goods using a scale of criteria. Autonomous Agents and Multi-Agent Systems **30**(2), 259–290 (2016)
19. Bredereck, R., Kaczmarczyk, A., Niedermeier, R.: Envy-free allocations respecting social networks. In: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-18), pp. 283–291. ACM, Stockholm, Sweden (2018)
20. Burkard, R., Dell'Amico, M., Martello, S.: Assignment Problems. Society for Industrial and Applied Mathematics (2012)
21. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M.: On low-envy truthful allocations. In: Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT-09), pp. 111–119. Springer, Venice, Italy (2009)
22. Cato, S.: Local strict envy-freeness in large economies. Mathematical Social Sciences **59**(3), 319–322 (2010)
23. Chevaleyre, Y., Endriss, U., Maudet, N.: Distributed fair allocation of indivisible goods. Artificial Intelligence **242**, 1–22 (2017)
24. Damamme, A., Beynier, A., Chevaleyre, Y., Maudet, N.: The power of swap deals in distributed resource allocation. In: Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-15), pp. 625–633. ACM, Istanbul, Turkey (2015)
25. Dickerson, J.P., Goldman, J.R., Karp, J., Procaccia, A.D., Sandholm, T.: The computational rise and fall of fairness. In: Proceedings of the 28th Conference on Artificial

Intelligence (AAAI-14), pp. 1405–1411. AAAI Press, Québec City, Québec, Canada (2014)
26. Feldman, A., Kirman, A.: Fairness and envy. The American Economic Review **64**(6), 995–1005 (1974)
27. Feldman, M., Lai, K., Zhang, L.: The proportional-share allocation market for computational resources. IEEE Transactions on Parallel and Distributed Systems **20**(8), 1075–1088 (2009)
28. Fellows, M.R., Hermelin, D., Rosamond, F., Vialette, S.: On the parameterized complexity of multiple-interval graph problems. Theoretical Computer Science **410**(1), 53–61 (2009)
29. Festinger, L.: A theory of social comparison process. Human Relations **7**(2), 117–140 (1954)
30. Flammini, M., Mauro, M., Tonelli, M.: On social envy-freeness in multi-unit markets. In: Proceedings of the 32nd Conference on Artificial Intelligence (AAAI-18), pp. 1031–1038. AAAI Press, New Orleans, Louisiana, USA (2018)
31. Foley, D.K.: Resource allocation and the public sector. Yale Economic Essays **7**(1), 45–98 (1967)
32. Garey, M.R., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-Completeness. W. H. Freeman & Co., New York, NY, USA (1979)
33. Gehrlein, W.V., Fishburn, P.C.: The probability of the paradox of voting: A computable solution. Journal of Economic Theory **13**(1), 14 – 25 (1976)
34. Gourvès, L., Lesca, J., Wilczynski, A.: Object allocation via swaps along a social network. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-17), pp. 213–219. International Joint Conferences on Artificial Intelligence, Melbourne, Australia (2017)
35. Hagberg, A.A., Schult, D.A., Swart, P.J.: Exploring network structure, dynamics, and function using NetworkX. In: Proceedings of the 7th Python in Science Conference (SciPy-08), pp. 11 – 15. Pasadena, CA USA (2008)
36. Hylland, A., Zeckhauser, R.: The efficient allocation of individuals to positions. Journal of Political Economy **87**(2), 293–314 (1979)
37. Igarashi, A., Peters, D.: Pareto-optimal allocation of indivisible goods with connectivity constraints. In: Proceedings of the 33rd Conference on Artificial Intelligence (AAAI-19). To appear, Honolulu, Hawaii, USA (2019)
38. de Keijzer, B., Bouveret, S., Klos, T., Zhang, Y.: On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences. In: Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT-09), pp. 98–110. Springer, Venice, Italy (2009)
39. Kleinberg, J.M., Tardos, É.: Algorithm Design. Addison-Wesley (2006)
40. Kondratev, A.Y., Nesterov, A.S.: Minimal envy and popular matchings. CoRR **abs/1902.08003** (2019)
41. Lipton, R.J., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of indivisible goods. In: Proceedings of the 5th ACM Conference on Electronic Commerce (EC-04), pp. 125–131. ACM, New York, NY, USA (2004)
42. Nguyen, T.T., Rothe, J.: Minimizing envy and maximizing average Nash social welfare in the allocation of indivisible goods. Discrete Applied Mathematics **179**, 54–68 (2014)
43. Paschos, V.T.: A ($\Delta/2$)-approximation algorithm for the maximum independent set problem. Information Processing Letters **44**(1), 11–13 (1992)
44. Saffidine, A., Wilczynski, A.: Constrained swap dynamics over a social network in distributed resource reallocation. In: Proceedings of the 11th International Symposium on Algorithmic Game Theory (SAGT-18), pp. 213–225. Springer, Beijing, China (2018)
45. Steger, A., Wormald, N.C.: Generating random regular graphs quickly. Combinatorics, Probability and Computing **8**(4), 377–396 (1999)
46. Suksompong, W.: Fairly allocating contiguous blocks of indivisible items. Discrete Applied Mathematics **260**, 227–236 (2019)
47. Svensson, L.G.: Strategy-proof allocation of indivisible goods. Social Choice and Welfare **16**(4), 557–567 (1999)
48. Vazirani, V.V.: Approximation Algorithms. Springer (2001)

49. Zhou, L.: On a conjecture by Gale about one-sided matching problems. Journal of Economic Theory **52**(1), 123 – 135 (1990)