The popular assignment problem: when cardinality is more important than popularity

Telikepalli Kavitha[1], Tamás Király[2], Jannik Matuschke[3], Ildikó Schlotter[4,5], and Ulrike Schmidt-Kraepelin[6]

[1] Tata Institute of Fundamental Research, Mumbai, India; kavitha@tifr.res.in
[2] Eötvös Loránd University, Budapest, Hungary; tamas.kiraly@ttk.elte.hu
[3] Research Center for Operations Management, KU Leuven, Leuven, Belgium; jannik.matuschke@kuleuven.be
[4] Budapest University of Technology and Economics, Hungary; ildi@cs.bme.hu
[5] Centre for Economic and Regional Studies, Budapest, Hungary;
[6] Technische Universität Berlin, Berlin, Germany; u.schmidt-kraepelin@tu-berlin.de

Abstract. We consider a matching problem in a bipartite graph $G = (A \cup B, E)$ where each node in A is an agent having preferences in partial order over her neighbors, while nodes in B are objects with no preferences. The size of our matching is more important than node preferences – thus, we are interested in maximum matchings only. Any pair of maximum matchings in G (equivalently, perfect matchings or *assignments*) can be compared by holding a head-to-head election between them where agents are voters. The goal is to compute an assignment M such that there is no *better* or "more popular" assignment. This is the popular assignment problem and it generalizes the well-studied popular matching problem (Abraham et al., 2007).

Popular assignments need not exist in every input instance. We show a polynomial-time algorithm that decides if the given instance admits one or not, and computes one, if so. In instances with no popular assignment, we consider the problem of finding an *almost* popular assignment, i.e., an assignment with minimum unpopularity margin. We show an $O^*(|E|^k)$ time algorithm for deciding if there exists an assignment with unpopularity margin at most k. We then show that this algorithm is essentially optimal by proving that the problem is NP-complete and $\mathsf{W}_l[1]$-hard with parameter k.

We also consider the minimum-cost popular assignment problem when there are edge costs, and show this problem to be NP-hard. This hardness holds even when all edge costs are in $\{0, 1\}$ and agents have strict preferences. By contrast, we propose a polynomial-time algorithm to the problem of deciding if there exists a popular assignment with a given set of forced/forbidden edges (this tractability holds even for partially ordered preferences). Our algorithms are combinatorial and based on LP duality. They search for an appropriate witness or *dual certificate*, and when a certificate cannot be found, we prove that the desired assignment does not exist in G.

1 Introduction

We consider a matching problem in a bipartite graph $G = (A \cup B, E)$ with *one-sided* preferences. Nodes in A, also called *agents*, have preferences in partial order over their neighbors while nodes in B, also called *objects*, have no preferences. This model is often called the *house allocation* problem as it arises in campus housing allocation in universities [1]. The fact that preferences are *one-sided* here makes this model very different from the *marriage* problem introduced by Gale and Shapley [13] in 1962, where *all* nodes have preferences over their neighbors.

Usually one seeks a matching in G that is *optimal* in some sense. Popularity is a well-studied notion of optimality in the model of one-sided preferences. Any pair of matchings, say M and N, can be compared by holding an election between them where agents are voters. Every agent prefers the matching where she gets assigned a more preferred partner and being unmatched is her worst choice. Let $\phi(M, N)$ be the number of agents who prefer M to N. Then we say that M is *more popular than* N if $\phi(M, N) > \phi(N, M)$. Let us write $\Delta(M, N) = \phi(M, N) - \phi(N, M)$.

Definition 1.1. *A matching M is* popular *if there is no matching more popular than M, i.e., $\Delta(M, N) \geq 0$ for all matchings N in G.*

The popular matching problem involves deciding if G admits a popular matching, and finding one if so. This is a well-studied problem from 2005, and there is an efficient algorithm to solve it [2].

Consider applications where the *size* of the matching is of primary importance. It is natural that as many students as possible be assigned campus housing. Another application is in assigning final year medical and nursing students to hospitals during emergencies (such as a pandemic) to overcome staff shortage [10]. Preferences of these students are important but the size of the matching is more important, since we want to augment human resources as much as possible. Thus what we seek is not a popular matching but a popular *maximum matching*, i.e., among maximum matchings, a best one. Our approach to prioritize the cardinality of the matching is in stark contrast with most existing results in the area of popular matchings, where the foremost requirement is usually popularity.

By augmenting G with dummy agents and artificial objects (see Section 3), we can assume that G admits a perfect matching, i.e., an *assignment*. So our problem becomes the popular perfect matching problem—we will call this the popular assignment problem in G. In other words, we seek an assignment of objects to agents such that every agent is assigned an object and, roughly speaking, there is no assignment that makes more agents happy (than it makes unhappy).

Definition 1.2. *A perfect matching M is a* popular assignment *if there is no perfect matching in G that is more popular than M, i.e., $\Delta(M, N) \geq 0$ for all perfect matchings N in G.*

Thus, a popular assignment is a *weak Condorcet winner* [6,24] where all perfect matchings are candidates and agents are voters. Weak Condorcet winners need not exist in a general voting instance; in our setting as well, a popular assignment need not exist in G. Consider the following simple example where $A = \{a_1, a_2, a_3\}$, $B = \{b_1, b_2, b_3\}$ and G is the complete bipartite graph $K_{3,3}$, i.e., every agent and object are adjacent. Suppose every agent has the same (strict) preference ordering: $b_1 \succ b_2 \succ b_3$, i.e., b_i is the i-th choice for $i = 1, 2, 3$. It is easy to check that for every assignment, there is a more popular assignment; so this instance has no popular assignment.

The popular assignment *problem.* Given a bipartite graph $G = (A \cup B, E)$ where every $a \in A$ has preferences in partial order over her neighbors, does G admit a popular assignment? If so, find one.

It is easy to show instances that admit popular assignments but do not have any popular matching (see Section 2.2). Interestingly, an algorithm for the popular assignment problem also solves the popular matching problem. By augmenting the given instance with artificial *worst-choice* objects and some dummy agents, we can construct an instance G' on at most twice as many nodes as in G such that G admits a popular matching if and only if G' admits a popular assignment (this simple reduction is given in Section 2.3). Thus, the popular assignment problem generalizes the popular matching problem.

By adjusting the usage of worst-choice objects appropriately, an algorithm for popular assignment can solve the following more general variant of both the popular matching problem and the popular assignment problem, and thus opens possibilities to a wide spectrum of applications.

Popularity with diversity. Consider instances $G = (A \cup B, E)$ where every agent has one of k colors associated with it, and we are interested in only those (not necessarily perfect) matchings that match for every $i \in \{1, \dots, k\}$, c_i agents of color i, where $s_i \leq c_i \leq t_i$ for some given bounds s_i and t_i, i.e., only those matchings that satisfy these lower and upper bounds for every color are admissible. We seek a matching that is popular *within* the set of admissible matchings (see Section 2.3 for a reduction to popular assignment).

Public housing programs constitute an application where such problems arise. For example, in Singapore, 70% of the residential real estate is managed by a public housing program which promotes ethnic diversity by imposing quotas on each housing block and ethnic group. Motivated by this market, Benabbou et al. [3] study a similar model with cardinal utilities.

Our contribution. Our first result is that the popular assignment problem can be solved in polynomial time. Let $|A| = |B| = n$ and $|E| = m$.

Theorem 1.1. *The* popular assignment *problem in* $G = (A \cup B, E)$ *can be solved in* $O(m \cdot n^{5/2})$ *time.*

When a popular assignment does not exist in G, a natural extension is to ask for an *almost* popular assignment, i.e., an assignment with *low* unpopularity. How do we measure the unpopularity of an assignment? A well-known measure is the *unpopularity margin* [21] defined for any assignment M as $\mu(M) = \max_N(\phi(N, M) - \phi(M, N)) = \max_N \Delta(N, M)$, where the maximum is taken over all assignments, that is, all perfect matchings N in G. Thus $\mu(M)$ is the maximum margin by which another assignment defeats M.

An assignment M is popular if and only if $\mu(M) = 0$. Let the k-unpopularity margin problem be the problem of deciding if G admits an assignment with unpopularity margin at most k. We generalize Theorem 1.1 to show the following result.

Theorem 1.2. *For any* $k \in \mathbb{Z}_{\geq 0}$*, the* k-unpopularity margin *problem in* $G = (A \cup B, E)$ *can be solved in* $O(m^{k+1} \cdot n^{5/2})$ *time.*

Rather than the exponential dependency on the parameter k in Theorem 1.2, can we solve the k-unpopularity margin problem in polynomial time? Or at least can we achieve a running time $f(k)\mathsf{poly}(m, n)$ for some function f so that the degree of the polynomial is independent of k? That is, can we get a *fixed-parameter tractable* algorithm with parameter k? The following hardness result shows that the algorithm of Theorem 1.2 is essentially optimal for the k-unpopularity margin problem. See Section 6.2 for the definition of $\mathsf{W}_l[1]$-hardness.

Theorem 1.3. *The* k-unpopularity margin *problem is* $\mathsf{W}_l[1]$*-hard with parameter* k *when agents' preferences are weak rankings, and it is* NP*-complete even if preferences are strict rankings.*

We next consider the minimum-cost popular assignment problem in G. So there is a cost function $c : E \to \mathbb{R}$ on the edges and a budget β and we want to know if G admits a popular assignment whose sum of edge costs is at most β. Computing a minimum-cost popular assignment efficiently would also imply an efficient algorithm for finding a popular assignment with *forced/forbidden* edges. We show the following hardness result.

Theorem 1.4. *The* minimum-cost popular assignment *problem is* NP*-complete, even if each edge cost is in* $\{0, 1\}$ *and agents have strict preferences.*

Interestingly, in spite of the above hardness result, the popular assignment problem with partial order preferences and forced/forbidden edges is tractable. Note that the assignment M must be popular among *all* assignments, not only those adhering to the forced and forbidden edge constraints. We show the following positive result.

Theorem 1.5. *Given a set $F^+ \subseteq E$ of forced edges and another set $F^- \subseteq E$ of forbidden edges, we can determine in polynomial time if there exists a popular assignment M in $G = (A \cup B, E)$ such that $F^+ \subseteq M$ and $F^- \cap M = \emptyset$.*

Thus the popular assignment problem is reminiscent of the well-known stable roommates problem[7]; in a roommates instance, finding a stable matching can be solved in polynomial time [15] even with forced/forbidden edges [12], however finding a minimum-cost stable matching is NP-hard [11].

1.1 Background

The notion of popularity in a marriage instance (where preferences are *two-sided* and strict) was introduced by Gärdenfors [14] in 1975. Popular matchings always exist in such an instance, since any stable matching is popular [14]. When preferences are one-sided, popular matchings need not exist. A simple and clean combinatorial characterization of popular matchings (see Section 2.1) was given in [2], leading to an $O(m\sqrt{n})$ time algorithm [2] for the popular matching problem. By contrast, a combinatorial characterization of popular assignments remains elusive. Finding a minimum unpopularity margin matching was proved to be NP-hard [21].

In the last fifteen years, popularity has been a widely studied concept. Researchers have considered extensions of the popular matching problem where one aims for a popular matching satisfying some additional optimality criteria such as rank-maximality or fairness [20,22], or where the notion of popularity is adjusted to incorporate capacitated objects or weighted agents [23,26]. Another variant of the popular matching problem was considered in [8] where nodes in A have strict preferences and nodes in B, i.e., objects, have no preferences, however each object cares to be matched to any of its neighbors. We refer to [7] for a survey on results in this area.

Among the literature on popular matchings, only a few studies have considered a setting that focuses on popularity within a restricted set of admissible solutions. The paper that comes closest to our work is [16] which considered the popular maximum matching problem in a marriage instance (where preferences are two-sided and strict). It was shown there that a popular maximum matching always exists in a marriage instance and one such matching can be computed in $O(mn)$ time. Very recently, it was shown in [18] that a minimum-cost popular maximum matching in a marriage instance can be computed in polynomial time. These results use the rich machinery of stable matchings in a marriage instance [13,25]. In contrast to these positive results for popular maximum matchings, computing an *almost*-stable maximum matching (one with the least number of blocking edges) in a marriage instance is NP-hard [4].

1.2 Techniques

Our popular assignment algorithm is based on LP duality. We show that a matching M is a popular assignment if and only if it has a *dual certificate* $\vec{\alpha} \in \{0, \pm 1, \ldots, \pm(n-1)\}^{2n}$ fulfilling certain constraints induced by the matching M. Our algorithm (see Section 4) can be viewed as a search for such a dual certificate. It associates a *level* $\ell(b)$ with every $b \in B$. This level function ℓ guides us in constructing a subgraph G_ℓ of G. If G_ℓ contains a perfect matching, then this matching is a popular assignment in G and the levels determine a corresponding dual certificate. If G_ℓ has no perfect matching, then we increase some ℓ-values and update G_ℓ accordingly, until eventually

[7] This problem asks for a stable matching in a general graph (which need not be bipartite) with strict preferences.

G_ℓ contains a perfect matching or the level of an object increases beyond the permitted range, in which case we prove that no popular assignment exists.

The LP method for popular matchings was introduced in [19] and dual certificates for popular matchings/maximum matchings in marriage instances were shown in [17,18]. However, dual certificates for popular matchings in instances with one-sided preferences have not been investigated so far. The existence of simple dual certificates for popular assignments is easy to show (see Section 3), but this does not automatically imply polynomial-time solvability. Our main novelty lies in showing a combinatorial algorithm to search for dual certificates in an instance G and in using this approach to solve the popular assignment problem in polynomial time.

Our other results. Our algorithm for the popular assignment problem with *forced/forbidden* edges (see Section 5) is a natural extension of the above algorithm where certain edges are excluded. The k-unpopularity margin algorithm (see Section 6) associates a *load* with every edge such that the total load is at most k and the overloaded edges are treated as *forced* edges. Our $W_l[1]$-hardness result shows that this $O^*(m^k)$ algorithm for the k-unpopularity margin problem is essentially optimal, i.e., it is highly unlikely that this problem admits an $f(k)m^{o(k)}$ algorithm for any computable function f. The NP-hardness for the minimum-cost popular assignment problem (see Section 7) implies that given a set of desired edges, it is NP-hard (even for strict preferences) to find a popular assignment that contains the maximum number of desired edges. Thus, although the *forced* edges variant is easy, the *desired* edges variant is hard.

2 Preliminaries

For any $v \in A \cup B$, let $\mathsf{Nbr}_G(v)$ denote the set of neighbors of v in G, and $\delta(v)$ the set of edges incident to v. For any $X \subseteq A \cup B$, we let $\mathsf{Nbr}_G(X) = \cup_{v \in X}\mathsf{Nbr}_G(v)$; we may omit the subscript G if it is clear from the context. For any set X of vertices (or edges) in G, let $G - X$ be the subgraph of G obtained by deleting the vertices (or edges, respectively) of X from G. For a matching M in G and a node v matched in M, we denote the partner of v by $M(v)$.

The preferences of node $a \in A$ on its neighbors are given by a strict partial order $\succ_a$, so $b \succ_a b'$ means that a prefers b to b'. We use $b \sim_a b'$ to denote that a is indifferent between b and b', i.e., neither $b \succ_a b'$ nor $b' \succ_a b$ holds. The relation $\succ_a$ is a *weak ranking* if $\sim_a$ is transitive. In this case, $\sim_a$ is an equivalence relation and there is a strict order on the equivalence classes. When each equivalence class has size 1, we call it a *strict ranking* or a *strict preference*.

2.1 A characterization of popular matchings from [2]

In order to characterize popular matchings, as done in [2], it will be convenient to add artificial worst-choice or last resort objects to the given instance $G = (A \cup B, E)$. So $B = B \cup \{l(a) : a \in A\}$, i.e., corresponding to each $a \in A$, a node $l(a)$ gets added to B and we set this node $l(a)$ as the worst-choice object for a. Thus we have $E = E \cup \{(a, l(a)) : a \in A\}$.

Let $E_1 = \{(a, b) \in E : b \text{ is a top-choice object for } a\}$. Call an object b *critical* if every maximum matching in $G_1 = (A \cup B, E_1)$ matches b, call b *non-critical* otherwise.

Theorem 2.1 ([2])**.** *A matching M in $G = (A \cup B, E)$ is popular if and only if M matches all critical objects and every agent a is matched to either one of her top-choice objects or one of her most preferred non-critical objects.*

2.2 An instance without popular matchings that admits a popular assignment

We describe a simple example that does not admit any popular matching, but admits a popular assignment. Let $G = (A \cup B, E)$ where $A = \{a_1, a_2, a_3\}$ and $B = \{b_1, b_2, b_3\}$ and the preference order of both a_1 and a_2 is $b_1 \succ b_2$ while the preference order of a_3 is $b_1 \succ b_2 \succ b_3$.

It follows from the characterization of popular matchings from [2] that a popular matching M has to match each of a_1, a_2, a_3 to either b_1 or b_2. Since this is not possible, this instance has no popular matching. It is easy to check that $M^* = \{(a_1, b_1), (a_2, b_2), (a_3, b_3)\}$ is a popular assignment in G.

2.3 Some simple reductions to the popular assignment problem

We will first show a reduction from the popular matching problem to the popular assignment problem. Let $G = (A \cup B, E)$ be an instance of the popular matching problem. Let $B' = B \cup \{l(a) : a \in A\}$. That is, corresponding to each $a \in A$, an object $l(a)$ (the last resort of a) is in B' and we set this object $l(a)$ as the worst-choice of a. Let $A' = A \cup \{d_1, \ldots, d_{|B|}\}$, i.e., there are $|B|$ many dummy agents in A'. Each dummy agent d_i is adjacent to all objects in B' and is indifferent between any two of them.

It is easy to see that every matching M in G can be extended to a *perfect matching* M' in this new graph $G' = (A' \cup B', E')$ and conversely, every perfect matching M' in G' projects to a matching M in G. For any pair of matchings M and N in G, observe that $\Delta(M, N) = \Delta(M', N')$. Thus an algorithm that finds a popular assignment in G' solves the popular matching problem in G.

Popularity with diversity. Recall this problem defined in Section 1 where every agent in an instance $G = (A \cup B, E)$ has one of k colors associated with it, and admissible matchings are those that for every $i \in \{1, \ldots, k\}$ match c_i agents of color i where $s_i \le c_i \le t_i$ for some given bounds s_i and t_i. We seek a matching that is popular within the set of admissible matchings.

We augment B by adding $n_i - s_i$ artificial objects for each i, where n_i is the number of agents colored i. For each i, these $n_i - s_i$ objects are tied as the worst-choices of all agents colored i. Let $A' = A \cup \{d_1, \ldots, d_{n'}\}$, where $n' = |B| - \sum_i s_i$. Every dummy agent $d \in \{d_1, \ldots, d_{n'}\}$ is adjacent to all objects in B and for each color i some fixed $t_i - s_i$ artificial objects meant for color class i introduced above—as before, d is indifferent between any two of its neighbors. So for each i, there are $n_i - t_i$ artificial objects not adjacent to any dummy agent. Let G' be the new instance. It is easy to see that an algorithm that finds a popular assignment in G' solves our problem in G.

3 Dual Certificates for Popular Assignments

Let $G = (A \cup B, E)$ be an input instance and let ν be the size of a maximum matching in G. Let us augment G with $|B| - \nu$ dummy agents that are adjacent to all objects in B (and indifferent among them), along with $|A| - \nu$ artificial objects that are tied as the worst-choice neighbors of all non-dummy agents. Any maximum matching M in the original graph extends to a perfect matching (i.e., assignment) M' in the augmented graph; moreover, $\Delta(M, N) = \Delta(M', N')$ for any pair of maximum matchings M and N in G. Thus, we can assume without loss of generality that the input instance G admits a perfect matching.

Let $|A| = |B| = n$ and $|E| = m$. Let M be any perfect matching in G. The following edge weight function wt_M in G will be useful. For any $(a, b) \in E$

$$\text{let } \mathsf{wt}_M(a,b) = \begin{cases} 1 & \text{if } a \text{ prefers } b \text{ to } M(a); \\ -1 & \text{if } a \text{ prefers } M(a) \text{ to } b; \\ 0 & \text{otherwise, i.e., if } b \sim_a M(a). \end{cases}$$

Let $\mathsf{wt}_M(N) = \sum_{e \in N} \mathsf{wt}_M(e)$ for any edge set $N \subseteq E$. Consider the following linear program LP1 and its dual LP2.

$$\max \sum_{e \in E} \mathsf{wt}_M(e) \cdot x_e \qquad \text{(LP1)}$$
$$\text{s.t.} \quad \sum_{e \in \delta(u)} x_e = 1 \quad \forall\, u \in A \cup B$$
$$x_e \geq 0 \quad \forall\, e \in E.$$

$$\min \sum_{u \in A \cup B} y_u \qquad \text{(LP2)}$$
$$\text{s.t.} \quad y_a + y_b \geq \mathsf{wt}_M(a,b) \quad \forall\, (a,b) \in E.$$

LP1 is well-known to be integral, and hence its optimal value is $\max_N \mathsf{wt}_M(N)$ where N is a perfect matching in G. The definition of wt_M implies that $\mathsf{wt}_M(N) = \Delta(N, M)$; recall that $\Delta(N, M) = \phi(N, M) - \phi(M, N)$. So M is a popular assignment if and only if the optimal value of LP1 is at most 0. In fact, the optimal value of LP1 is then exactly 0, by $\Delta(M, M) = 0$. Hence for a popular assignment M, the edge incidence vector of M is an optimal solution to LP1.

Theorem 3.1 gives a useful characterization of popular assignments. The proof of Theorem 3.1 is given in Section 6 along with the proof of a related result (Theorem 6.1). A *dual certificate* of a popular assignment M is an optimal solution $\vec{\alpha}$ to LP2 as given in Theorem 3.1.

Theorem 3.1. *M is a popular assignment if and only if there exists an optimal solution $\vec{\alpha}$ to LP2 such that $\alpha_a \in \{0, 1, 2, \ldots, (n-1)\}$ for all $a \in A$, $\alpha_b \in \{0, -1, -2, \ldots, -(n-1)\}$ for all $b \in B$, and $\sum_{u \in A \cup B} \alpha_u = 0$.*

4 The Popular Assignment Algorithm

The goal of our algorithm is to construct a perfect matching M in G along with a dual certificate $\vec{\alpha}$. Every $b \in B$ will have an associated *level* $\ell(b)$ in this algorithm and the α-value of b will be $-\ell(b)$, i.e., we set $\alpha_b = -\ell(b)$.

Given a function $\ell : B \to \mathbb{N}$ called a *level function*, for any $a \in A$ let $\ell^*(a) = \max_{b \in \mathsf{Nbr}(a)} \ell(b)$ be the highest level at which agent a has neighbors. Now we define the subgraph $G_\ell = (A \cup B, E_\ell)$ *induced by levels* $\ell(\cdot)$ by putting an edge $(a, b) \in E$ into E_ℓ if and only if

(i) b has level $\ell^*(a)$, and a has no neighbor in level $\ell^*(a)$ that she prefers to b, or
(ii) b has level $\ell^*(a) - 1$, and a prefers b to each of her neighbors in level $\ell^*(a)$, and moreover, a prefers none of her neighbors in level $\ell^*(a) - 1$ to b.

Thus in the subgraph G_ℓ, every agent has edges to her favorite *highest-level* neighbors and to her favorite neighbors one level below, provided these neighbors are preferred to all of her highest-level neighbors (see Fig. 1 for an illustration). The following lemma will be very useful.

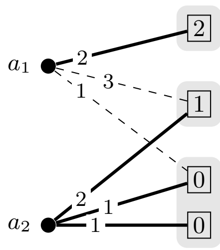

Fig. 1. Illustration of the subgraph G_ℓ of G for an instance with weak rankings and a level function ℓ. Circles indicate agents and squares objects; the ℓ-level of each object is written inside the square depicting it. Numbers on the edges indicate the agents' weak rankings. Bold edges are included in G_ℓ and dashed edges are not. All but two agents were omitted.

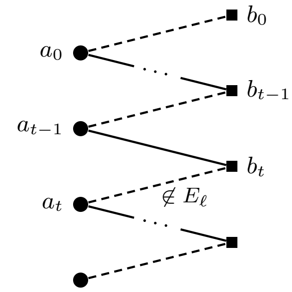

Fig. 2. An illustration of the M-augmenting path P within the proof of Lemma 4.2. Solid edges are in M and dashed edges are in $M^\star$. The edge (a_t, b_t) is not contained in E_ℓ.

Lemma 4.1. *A matching M in G is a popular assignment if and only if there exists a level function ℓ such that M is a perfect matching in G_ℓ. Further, this happens if and only if there is a level function ℓ and a dual certificate $\vec{\alpha}$ for M where $\ell(b) = |\alpha_b|$ for all $b \in B$ and M is a perfect matching in G_ℓ.*

Proof. Let us first show that if there exists a level function ℓ such that M is a perfect matching in G_ℓ, then M is a popular assignment in G. We construct a dual certificate for M as follows. Let $\alpha_b = -\ell(b)$ for all $b \in B$ and $\alpha_a = \ell(M(a))$ for all $a \in A$. Note that $\sum_{v \in A \cup B} \alpha_v = \sum_{a \in A}(\alpha_a + \alpha_{M(a)}) = 0 = \sum_{e \in M} \mathsf{wt}_M(e)$. Thus $\vec{\alpha}$ is optimal for LP2 if it is feasible. It remains to show that $\alpha_a + \alpha_b \geq \mathsf{wt}_M(a, b)$ for every $(a, b) \in E$.

So let $(a, b) \in E$ and let $b' := M(a)$. Note that $\alpha_a + \alpha_b = \ell(b') - \ell(b)$. We show that $\ell(b') - \ell(b) \geq \mathsf{wt}_M(a, b)$. Because $(a, b') \in E_\ell$, one of the following cases holds:

- Case (i): $\ell(b') = \ell^*(a)$, so a prefers no neighbor of hers in level $\ell^*(a)$ to b'. We have two subcases:
  - If $\ell(b) = \ell^*(a)$, then a does not prefer b to b' and hence $\mathsf{wt}_M(a, b) \leq 0 = \ell(b') - \ell(b)$.
  - If $\ell(b) < \ell^*(a)$, then $\ell(b') - \ell(b) \geq 1 \geq \mathsf{wt}_M(a, b)$.
- Case (ii): $\ell(b') = \ell^*(a) - 1$, so a prefers b' to each of her neighbors in level $\ell^*(a)$, and a prefers none of her neighbors in level $\ell^*(a) - 1$ to b'. We have three subcases:
  - If $\ell(b) = \ell^*(a)$, then a prefers b' to b and hence $\mathsf{wt}_M(a, b) = -1 = \ell(b') - \ell(b)$.
  - If $\ell(b) = \ell^*(a) - 1$, then a does not prefer b to b' and hence $\mathsf{wt}_M(a, b) \leq 0 = \ell(b') - \ell(b)$.
  - If $\ell(b) < \ell^*(a) - 1$, then $\ell(b') - \ell(b) \geq 1 \geq \mathsf{wt}_M(a, b)$.

Thus in each of these cases $\mathsf{wt}_M(a, b) \leq \ell(b') - \ell(b) = \alpha_a + \alpha_b$. Hence α is a dual certificate for M, and thus M is a popular assignment by Theorem 3.1.

We will now show the converse. Let M be a popular assignment in G and let $\vec{\alpha}$ be a dual certificate for M. We claim that M is a matching in the graph G_{ℓ_α} induced by levels $\ell_{\vec{\alpha}}$ with $\ell_{\vec{\alpha}}(b) = |\alpha_b|$ for all $b \in B$. To prove this, we use that $\alpha_a + \alpha_b \geq \mathsf{wt}_M(a, b)$ for every $(a, b) \in E$. First, because the incidence vector of M and $\vec{\alpha}$ are optimal solutions to LP1 and LP2, respectively, we get $\alpha_a + \alpha_{M(a)} = 0$ for each $a \in A$ by complementary slackness. This implies $\alpha_a = -\alpha_{M(a)} = \ell_{\vec{\alpha}}(M(a))$. Therefore, $\ell_{\vec{\alpha}}(M(a)) \geq \ell_{\vec{\alpha}}(b) + \mathsf{wt}_M(a, b)$ for all $(a, b) \in E$.

Since $\mathsf{wt}_M(e) \geq -1$ for all edges e, any agent a has to be matched in M to either (i) an *undominated* neighbor in level $\ell^*_{\vec{\alpha}}(a)$ (i.e., a prefers none of her neighbors in this level to $M(a)$) or

(ii) an undominated neighbor in level $\ell^*_{\vec{\alpha}}(a) - 1$ which, moreover, has to *dominate* (i.e., be preferred by a to) all of a's neighbors in level $\ell^*_{\vec{\alpha}}(a)$. So M is a perfect matching in $G_{\ell_{\vec{\alpha}}}$. □

The algorithm. Consider Algorithm 1 on input $G = (A \cup B, E)$. In the search for a dual certificate, this algorithm will maintain a *level* $\ell(b)$ for every $b \in B$. Initially, $\ell(b) = 0$ for every $b \in B$.

Our algorithm checks whether there exists a popular assignment by computing a perfect matching in the graph G_ℓ. If no such matching exists, the levels of unmatched objects are increased, the graph G_ℓ is updated accordingly, and the search continues.

Eventually, either a perfect matching in G_ℓ is found, or the level of an object exceeds $n-1$. In the latter case we can conclude that no popular assignment exists, as we will show below.

Algorithm 1 Finding a popular assignment

for all $b \in B$ **do** $\ell(b) = 0$.
while $\ell(b) < n$ for all $b \in B$ **do**
  Construct the graph G_ℓ and compute a maximum matching M in G_ℓ.
  if M is a perfect matching **then return** M.
  for all $b \in B$ unmatched in M **do** $\ell(b) = \ell(b) + 1$.
return "G has no popular assignment".

Running time. Computing a maximum matching in G_ℓ takes $O(m\sqrt{n})$ time. In every iteration of the algorithm, the value $\sum_{b \in B} \ell(b)$ increases. So the number of iterations is at most n^2. Hence the running time of our algorithm is $O(m \cdot n^{5/2})$.

Theorem 4.1. *If our algorithm returns a matching M, then M is a popular assignment in G.*

Theorem 4.1 follows immediately from Lemma 4.1. The more difficult part in our proof of correctness is to show that whenever our algorithm says that G has no popular assignment, the instance G indeed has no popular assignment. This is implied by Theorem 4.2.

Theorem 4.2. *Let $M^\star$ be a popular assignment in G and let $\vec{\alpha}$ be a dual certificate of $M^\star$. Then for every $b \in B$, we have $|\alpha_b| \geq \ell(b)$, where $\ell(b)$ is the level of b when our algorithm terminates.*

If our algorithm terminates because $\ell(b) = n$ for some $b \in B$, then $|\alpha_b| \geq n$ for any dual certificate $\vec{\alpha}$ by Theorem 4.2. However $|\alpha_b| \leq n-1$ by definition, a contradiction. So G has no popular assignment.

The following lemma is crucial for proving Theorem 4.2. It guarantees that when the algorithm increases $\ell(b)$ for some unmatched object $b \in B$, then the new level function does not exceed $|\alpha_b|$.

Lemma 4.2. *Let $M^\star$ be a popular assignment, let $\vec{\alpha}$ be a dual certificate of $M^\star$, and let $\ell : B \to \mathbb{N}$ be such that $\ell(b) \leq |\alpha_b|$ for all $b \in B$. Let M be a maximum matching in G_ℓ and let $b_0 \in B$ be an object that is left unmatched in M. Then $\ell(b_0) < |\alpha_{b_0}|$.*

Before we turn to the proof of Lemma 4.2, we point out that Theorem 4.2 follows from this lemma by a simple induction.

Proof (of Theorem 4.2). Let ℓ_i be the level function on the set B at the start of the i-th iteration of our algorithm. We are going to show by induction that for every i we have $|\alpha_b| \geq \ell_i(b)$ for all $b \in B$. This is true for $i = 1$, since $\ell_1(b) = 0$ for all $b \in B$. Now suppose that $|\alpha_b| \geq \ell_i(b)$ for all $b \in B$. Let $b_0 \in B$. If b_0 is matched in the maximum matching M in G_{ℓ_i}, then we know $\ell_{i+1}(b_0) = \ell_i(b_0) \leq |\alpha_{b_0}|$. If b_0 is left unmatched in M, then $\ell_{i+1}(b_0) = \ell_i(b_0) + 1$ and $\ell_i(b_0) < |\alpha_{b_0}|$ by Lemma 4.2. Thus $\ell_{i+1}(b_0) \leq |\alpha_{b_0}|$ in either case, completing the induction. □

Proof (of Lemma 4.2). By Lemma 4.1, $M^\star$ is a perfect matching in $G_{\ell_{\vec{\alpha}}}$ where $\ell_{\vec{\alpha}}(b) = |\alpha_b|$ for $b \in B$. Thus, the symmetric difference $M \oplus M^\star$ in G contains an M-augmenting path P starting at b_0. However, as M is of maximum size in G_ℓ, the path P must contain an edge that is not in E_ℓ.

Let $(b_0, a_0, b_1, a_1, \ldots, b_t, a_t)$ be any prefix of P such that $(a_t, b_t) \notin E_\ell$ (see Fig. 2). Note that (a_0, b_0) and (a_t, b_t) are in $M^\star$, since M leaves b_0 unmatched and $M \subseteq E_\ell$. Thus $(a_h, b_h) \in M^\star \subseteq E_{\ell_{\vec{\alpha}}}$ for all $h \in \{0, \ldots, t\}$ and $(a_h, b_{h+1}) \in M \subseteq E_\ell$ for all $h \in \{0, \ldots, t-1\}$. We will show that $\ell(b_h) < \ell_{\vec{\alpha}}(b_h)$ for all $h \in \{0, \ldots, t\}$, and thus in particular, $\ell(b_0) < \ell_{\vec{\alpha}}(b_0) = |\alpha_{b_0}|$.

We first show that $\ell(b_t) < \ell_{\vec{\alpha}}(b_t)$. Assume for contradiction that $\ell(b_t) = \ell_{\vec{\alpha}}(b_t)$. Using the fact that $(a_t, b_t) \notin E_\ell$, one of the following cases must hold:

- a_t has a neighbor in level at least $\ell(b_t) + 2$, or
- a_t has a neighbor in level $\ell(b_t) + 1$ that is not dominated by b_t, or
- a_t has a neighbor in level $\ell(b_t)$ that is preferred to b_t.

As $\ell_{\vec{\alpha}}(b_t) = \ell(b_t)$ and $\ell_{\vec{\alpha}}(b) \geq \ell(b)$ for all $b \in B$, in each case we get $(a_t, b_t) \notin E_{\ell_{\vec{\alpha}}}$, a contradiction.

Now suppose there is $h \in \{0, \ldots, t-1\}$ with $\ell(b_{h+1}) < \ell_{\vec{\alpha}}(b_{h+1})$ but $\ell(b_h) = \ell_{\vec{\alpha}}(b_h)$. Recall that $(a_h, b_{h+1}) \in M \subseteq E_\ell$, which leaves us with the following possibilities:

- $\ell(b_{h+1}) \geq \ell(b_h) + 1$: then $\ell_{\vec{\alpha}}(b_{h+1}) \geq \ell_{\vec{\alpha}}(b_h) + 2$;
- $\ell(b_{h+1}) = \ell(b_h)$: then a_h does not prefer b_h to b_{h+1}, but $\ell_{\vec{\alpha}}(b_{h+1}) \geq \ell_{\vec{\alpha}}(b_h) + 1$;
- $\ell(b_{h+1}) = \ell(b_h) - 1$: then a_h prefers b_{h+1} to b_h, but $\ell_{\vec{\alpha}}(b_{h+1}) \geq \ell_{\vec{\alpha}}(b_h)$.

In each of these cases, we get $(a_h, b_h) \notin E_{\ell_{\vec{\alpha}}}$ by the definition of $G_{\ell_{\vec{\alpha}}}$, again a contradiction. □

We remark that if a popular assignment exists, then the algorithm returns a popular assignment M and a corresponding dual certificate $\vec{\alpha}$ such that $\ell_{\vec{\alpha}} \leq \ell_{\vec{\alpha}'}$ for any dual certificate $\vec{\alpha}'$ of any popular assignment M'. This shows that there is a unique minimal dual certificate in this sense.

We close this section by pointing out a generalization of Lemma 4.2 that encapsulates the main argument of the preceding proof. This insight will be useful for generalizing our algorithmic result in the next two sections.

Lemma 4.3. *Let $\ell, \ell' : B \to \mathbb{N}$ be such that $\ell(b) \leq \ell'(b)$ for all $b \in B$. Let M and M' be matchings in G_ℓ and $G_{\ell'}$, respectively. Let $b_0 \in B$ be an object that is matched in M' but not in M. Let P be the path in $M \oplus M'$ containing b_0. If P contains an edge not in E_ℓ, then $\ell(b_0) < \ell'(b_0)$.*

5 Finding a Popular Assignment with Forced/Forbidden Edges

In this section we consider a variant of the popular assignment problem where, in addition to our instance, we are given a set F^+ of *forced* edges and a set F^- of *forbidden* edges, and we are looking for a popular assignment that contains F^+ and is disjoint from F^-. Observe that it is sufficient to deal with forbidden edges, since putting an edge (a, b) into F^+ is the same as putting all the edges in the set $\{(a, b') : b' \in \mathsf{Nbr}(a) \text{ and } b' \neq b\}$ into F^-.

The popular assignment with forbidden edges *problem.* Given a bipartite graph $G = (A \cup B, E)$ where every $a \in A$ has preferences in partial order over her neighbors, together with a set $F \subseteq E$ of forbidden edges, does G admit a popular assignment M avoiding F, i.e., one where $M \cap F = \emptyset$?

We will show that in order to deal with forbidden edges, it suffices to modify our algorithm in Section 4 as follows; see Algorithm 2. The only difference from the earlier algorithm is that we find a maximum matching in the subgraph $G_\ell - F$, i.e., on the edge set $E_\ell \setminus F$.

Algorithm 2 Finding a popular assignment with forbidden edges

1: **for all** $b \in B$ **do** $\ell(b) = 0$.
2: **while** $\ell(b) < n$ for all $b \in B$ **do**
3: Construct the graph $G_\ell = (A \cup B, E_\ell)$ and find a maximum matching M in $G_\ell - F$.
4: **if** M is a perfect matching **then return** M.
5: **for all** $b \in B$ unmatched in M **do** $\ell(b) = \ell(b) + 1$.
6: **return** "G has no popular assignment with forbidden set F".

Theorem 5.1. *The above algorithm outputs a popular assignment avoiding F, if such an assignment exists in G.*

Proof. Recall that any perfect matching in G_ℓ is a popular assignment in G by Lemma 4.1. It is therefore immediate that if the above algorithm outputs a matching M, then M is a popular assignment in G that avoids F.

Let us now prove that if there exists a popular assignment $M^\star$ avoiding F, then our algorithm outputs such an assignment. Let $\vec{\alpha}$ be a dual certificate for $M^\star$. Let ℓ_i denote the level function at the beginning of iteration i of the algorithm. As in the proof of Theorem 4.2, we show by induction that $\ell_i(b) \le |\alpha_b|$ for all $b \in B$ for any iteration i.

This is clearly true initially with $\ell_1(b) = 0$ for all $b \in B$. To complete the induction, it suffices to show that $\ell_i(b_0) < |\alpha_{b_0}|$ for all $b_0 \in B$ that are unmatched by any maximum matching M in G_{ℓ_i}. Since $M^\star$ is a perfect matching, the symmetric difference $M \oplus M^\star$ contains an M-augmenting path P starting at b_0. However, because M has maximum size in $G_{\ell_i} - F$, the path P must contain an edge $e \notin E_{\ell_i} \setminus F$. We have $(M \cup M^\star) \cap F = \emptyset$, thus we obtain $e \notin F$ and therefore $e \notin E_{\ell_i}$. Note that Lemma 4.1 implies $M^\star \subseteq E_{\ell_{\vec{\alpha}}}$ (recall that $\ell_{\vec{\alpha}}(b) = |\alpha_b|$ for all $b \in B$). Furthermore, $\ell_i(b) \le \ell_{\vec{\alpha}}(b)$ for all $b \in B$ by our induction hypothesis. We can thus apply Lemma 4.3 with $M' = M^\star$ and $\ell' = \ell_{\vec{\alpha}}$ to obtain $\ell_i(b_0) < \ell_{\vec{\alpha}}(b_0) = |\alpha_{b_0}|$, which completes the induction step. ◻

6 Finding an Assignment with Minimum Unpopularity Margin

In this section we consider the k-unpopularity-margin problem in G. Section 6.1 has our algorithmic result and Section 6.2 and Section 6.3 contain our hardness results.

6.1 Our algorithm

For any assignment M, recall that the optimal value of LP1 is $\max_N \Delta(N, M) = \mu(M)$, where the maximum is taken over all assignments N in G. Consequently, $\mu(M)$ equals the optimal value of the dual linear program LP2 as well. Therefore, $\mu(M) = k$ if and only if there exists an optimal solution $\vec{\alpha}$ to LP2 for which $\sum_{u \in A \cup B} \alpha_u = k$. This leads us to a characterization of assignments with a bounded unpopularity margin that is a direct analog of Theorem 3.1.

Theorem 6.1. *M is an assignment with $\mu(M) \leq k$ if and only if there exists a solution $\vec{\alpha}$ to LP2 such that $\alpha_a \in \{0, 1, \ldots, n\}$ for all $a \in A$, $\alpha_b \in \{0, -1, \ldots, -(n-1)\}$ for all $b \in B$, and $\sum_{u \in A \cup B} \alpha_u \leq k$.*

Proof. (Proof of Theorems 3.1 and 6.1) If there exists an optimal solution $\vec{\alpha}$ to LP2 such that $\sum_{u \in A \cup B} \alpha_u = 0$, then the optimal value of LP2 is 0 and hence by LP-duality, the optimal value of LP1 is also 0. Thus $\Delta(N, M) \leq 0$ for any perfect matching N; in other words, M is a popular assignment. Similarly, if there is an optimal solution $\vec{\alpha}$ to LP2 such that $\sum_{u \in A \cup B} \alpha_u \leq k$, then $\Delta(N, M) \leq k$ for any perfect matching N, so $\mu(M) \leq k$.

We will now show the converse. Let M be a perfect matching with $\mu(M) = k$; there exists a dual optimal solution $\vec{\alpha}$ such that $\sum_{u \in A \cup B} \alpha_u = k$. Moreover, we can assume $\vec{\alpha} \in \mathbb{Z}^{2n}$ due to the total unimodularity of the constraint matrix. We can also assume $\alpha_b \leq 0$ for all $b \in B$, because feasibility and optimality are preserved if we decrease α_b for all $b \in B$ and increase α_a for all $a \in A$ by the same amount.

Let us choose $\vec{\alpha}$ such that $\sum_{b \in B} \alpha_b$ is maximal subject to the above assumptions. We claim that if there is no $b \in B$ with $\alpha_b = -r$ for some $r \in \mathbb{N}$, then there is no $b \in B$ with $\alpha_b \leq -(r+1)$. Suppose the contrary, and let $B' = \{b \in B : \alpha_b < -r\}$ and $A' = \{M(b) : b \in B'\}$. Since $\alpha_a + \alpha_b \geq 0$ for every $(a, b) \in M$, we have $\alpha_a \geq r + 1$ for every $a \in A'$. Let $\vec{\alpha}'$ be obtained by decreasing α_a by 1 for all $a \in A'$ and increasing α_b by 1 for all $b \in B'$. The dual feasibility constraints $\alpha'_a + \alpha'_b \geq \mathsf{wt}_M(a, b)$ can only be violated if $a \in A'$, $b \notin B'$, and $\alpha_a + \alpha_b = \mathsf{wt}_M(a, b)$. But this would imply $\alpha_a \geq r + 1$, $\alpha_b \geq -r + 1$ (since $b \notin B'$ and α_b cannot be $-r$), and $\alpha_a + \alpha_b = \mathsf{wt}_M(a, b) \leq 1$, a contradiction. Thus, $\vec{\alpha}'$ is also an optimal dual solution, and $\sum_{b \in B} \alpha'_b > \sum_{b \in B} \alpha_b$, contradicting the choice of $\vec{\alpha}$.

We have shown that the values that $\vec{\alpha}$ takes on B are consecutive integers, so we obtain that $\alpha_b \in \{0, -1, -2, \ldots, -(n-1)\}$ for all $b \in B$. Since $\alpha_a + \alpha_b \geq 0$ for every $(a, b) \in M$, we have $\alpha_a \geq 0$ for every $a \in A'$.

To conclude the proof of Theorem 3.1, observe that M is an optimal primal solution when $k = 0$, so $\alpha_a + \alpha_b = 0$ for every $(a, b) \in M$. This implies that $\alpha_a \in \{0, 1, 2, \ldots, n-1\}$ for all $a \in A$. As for Theorem 6.1, let N be a perfect matching that is optimal for LP1; then $\alpha_a + \alpha_b = \mathsf{wt}_M(a, b) \leq 1$ for every $(a, b) \in N$ by complementary slackness, and therefore $\alpha_a \leq n$ for all $a \in A$. $\square$

Generalizing the notion that we already used for popular assignments, we define a *dual certificate* for an assignment M with unpopularity margin k as a solution $\vec{\alpha}$ to LP2 with properties as described in Theorem 6.1. So let us suppose that M is an assignment with $\mu(M) = k$ and $\vec{\alpha}$ is a dual certificate for M. We define the *load* of $(a, b) \in M$ as $\alpha_a + \alpha_b$, and we will say that an edge $(a, b) \in M$ is *overloaded* (with respect to $\vec{\alpha}$), if it has a positive load, that is, $\alpha_a + \alpha_b > 0$. Clearly, the total load of all edges in M is at most k, moreover $\alpha_a + \alpha_b \geq \mathsf{wt}_M(a, b) = 0$ for every $(a, b) \in M$, so there are at most k overloaded edges in M.

Given a level function $\ell : B \to \mathbb{N}$ and an integer $\lambda \in \mathbb{N}$, we say that edge (a, b) is λ-*feasible*, if

(i) b has level at least $\ell^*(a) - \lambda + 1$ where $\ell^*(a) = \max_{b \in \mathsf{Nbr}(a)} \ell(b)$, or
(ii) b has level $\ell^*(a) - \lambda$ and a has no neighbor in level $\ell^*(a)$ that she prefers to b, or
(iii) b has level $\ell^*(a) - \lambda - 1$, a prefers b to each of her neighbors in level $\ell^*(a)$, and moreover, a prefers none of her neighbors in level $\ell^*(a) - 1$ to b.

Note that 0-feasible edges are exactly those contained in the graph G_ℓ induced by levels $\ell(\cdot)$, as defined in Section 4. The following observation follows directly from the constraints of LP2.

Proposition 6.1. *Consider the level function $\ell_{\vec{\alpha}}$ where the level of any $b \in B$ is $\ell_{\vec{\alpha}}(b) = -\alpha_b$. Then any edge $e \in M$ with load λ is λ-feasible.*

Given a level function $\ell : B \to \mathbb{N}$ and a load capacity function $\lambda : E \to \mathbb{N}$, we define the graph $G_{\ell,\lambda} = (A \cup B, E_{\ell,\lambda})$ *induced by levels* $\ell(\cdot)$ *and load capacities* $\lambda(\cdot)$ by putting an edge e into $E_{\ell,\lambda}$ if and only if e is $\lambda(e)$-feasible.

We are now ready to describe an algorithm for finding an assignment M with $\mu(M) \leq k$ if such an assignment exists. Algorithm 3 starts by guessing the load $\lambda(e)$ for each edge e of E. Then we use a variant of the algorithm for Theorem 5.1 that enables each edge e with $\lambda(e) > 0$ to have positive load (so $G_{\ell,\lambda}$ will be used instead of G_ℓ), and treats the overloaded edges as forced edges.

Algorithm 3 Finding a popular assignment with unpopularity margin at most k

1: **for all** functions $\lambda : E \to \mathbb{N}$ with $\sum_{e \in E} \lambda(e) \leq k$ **do**
2:     Set $K = \{e \in E : \lambda(e) > 0\}$ as the edges which we will overload.
3:     Set $F = \{(a, b') \in E : (a, b) \in K, b' \neq b\}$ as the set of forbidden edges.
4:     **for all** $b \in B$ **do** $\ell(b) = 0$.
5:     **while** $\ell(b) < n$ for all $b \in B$ **do**
6:         Construct the graph $G_{\ell,\lambda}$ and find a maximum matching M in $G_{\ell,\lambda} - F$.
7:         **if** M is a perfect matching **then return** M.
8:         **for all** $b \in B$ unmatched in M **do** $\ell(b) = \ell(b) + 1$.
9: **return** "G has no assignment M with $\mu(M) \leq k$".

Observe that there are at most m^k ways to choose the load capacity function λ, where $m = |E|$, by the bound $\sum_{e \in E} \lambda(e) \leq k$. Each iteration of the while-loop takes $O(m\sqrt{n})$ time and there are at most $m^k \cdot n^2$ such iterations. Thus the running time of the above algorithm is $O(m^{k+1} \cdot n^{5/2})$.

Proof of Theorem 1.2. First, suppose that Algorithm 3 outputs an assignment M. Consider the values of $\lambda(\cdot)$ and $\ell(\cdot)$ at the moment the algorithm outputs M. Set $\alpha_b = -\ell(b)$ for each object, and $\alpha_a = -\alpha_b + \lambda(a, b)$ for each edge $(a, b) \in M$. From the definition of $G_{\ell,\lambda}$ and λ-feasibility, such a vector $\vec{\alpha}$ fulfills all constraints in LP2. Hence, by $\sum_{u \in A \cup B} \alpha_u = \sum_{(a,b) \in M} \lambda(a, b) \leq k$, we get that $\mu(M) \leq k$.

Second, assume that G admits an assignment $M^\star$ with $\mu(M^\star) \leq k$, and let $\vec{\alpha}$ be a dual certificate for $M^\star$. We need to show that our algorithm will produce an output. Consider those iterations where $\lambda(e)$ equals the load of each edge $e \in M^\star$; we call this the *significant branch* of the algorithm. We claim that $|\alpha_b| \geq \ell(b)$ holds throughout the run of the significant branch.

To prove our claim, we use the same approach as in the proof of Theorem 5.1, based on induction. Clearly, the claim holds at the beginning of the branch; we need to show that $\ell(b)$ is increased only if $|\alpha_b| > \ell(b)$. So let $|\alpha_b| \geq \ell(b)$ for each $b \in B$ at the beginning of an iteration (steps (6)–(8)), and consider an object b whose value is increased at the end of the iteration.

First, assume that b is incident to some overloaded edge (a, b) of $M^\star$ with load $\lambda(a, b)$. Since b is not matched in a maximum matching in $G_{\ell,\lambda} - F$, we know that the edge $(a, b) \in K$ is not λ-feasible with respect to $\ell(\cdot)$. However, by Proposition 6.1, (a, b) is $\lambda(a, b)$-feasible with respect to the level function $\ell_{\vec{\alpha}}$ where $\ell_{\vec{\alpha}}(b) = |\alpha_b|$. Recall that by our induction hypothesis, the $\ell_{\vec{\alpha}}$-level of any object is at least its ℓ-level. Moreover, increasing the level of any object other than b cannot result in making the edge (a, b) a $\lambda(a, b)$-feasible edge. It follows that $\ell_{\vec{\alpha}}(b) = |\alpha_b| > \ell(b)$ must hold.

Second, assume that b is not incident to any overloaded edge. Consider the path P in $M \oplus M^\star$ on which b lies. Notice that, since the edges of F are treated as forbidden edges, M assigns each agent incident to some edge in K either the object assigned to it by $M^\star$, or does not assign any object to it. Therefore, no such agent lies on the path P. Consequently, we can apply Lemma 4.3 to the matchings $M \cap P \subseteq E_\ell$ and $M^\star \cap P \subseteq E_{\ell_{\vec{\alpha}}}$. This yields $|\alpha_b| = \ell_{\vec{\alpha}}(b) > \ell(b)$, proving our claim.

Finally, note that in the significant branch, no object may have ℓ-level higher than $n - 1$, as implied by the properties of $\vec{\alpha}$ stated in Theorem 6.1. $\square$

6.2 Hardness results

We now contrast Theorem 1.2 with Theorem 1.3 which states that finding an assignment with minimum unpopularity margin is NP-hard and $\mathsf{W}_l[1]$-hard with respect to the parameter k, our bound on the unpopularity margin. A parameterized problem Q is $\mathsf{W}_l[1]$-hard if there exists a linear FPT-reduction from the weighted antimonotone cnf 2-sat (or wcnf 2sat$^-$) problem[8] to Q [5]. Since wcnf 2sat$^-$ is a W[1]-complete problem [9], $\mathsf{W}_l[1]$-hardness implies W[1]-hardness. While W[1]-hardness of k-unpopularity margin shows that it cannot be solved in time $f(k)|I|^{O(1)}$ for any computable function f unless $W[1] = FPT$ (where $|I|$ denotes the input length), the results of Chen et al. [5] enable us to obtain a tighter lower bound: the $\mathsf{W}_l[1]$-hardness of k-unpopularity margin implies that it cannot even be solved in $f(k)|I|^{o(k)}$ time for any computable function f, unless all problems in SNP are solvable in subexponential time – a possibility widely considered unlikely. Therefore, Theorem 1.3 shows that our algorithm for Theorem 1.2 is essentially optimal.

Note that the unpopularity margin of any assignment M can be computed efficiently by determining the optimal value of LP1, so the k-unpopularity margin problem is in NP. In the remainder of this section, we present a linear FPT-reduction from the Clique problem to the k-unpopularity margin problem where agents' preferences are weak rankings. By the work of Chen et al. [5], the $\mathsf{W}_l[1]$-hardness of k-unpopularity margin follows. Our reduction is a polynomial-time reduction as well, implying NP-hardness for the case of weak rankings; note that this also follows easily from the NP-hardness of finding a matching (not necessarily an assignment) with minimum unpopularity margin [21], using our reduction from the popular matching problem to the popular assignment problem. Both the reduction presented in this section and the reduction from popular matching in Section 2.3 use weak rankings. However, we prove the NP-hardness of k-unpopularity margin for strict rankings by reducing the problem with weak rankings to the case with strict rankings in Lemma 6.7 in Section 6.3.[9]

Instead of giving a direct reduction from Clique, we will use an intermediate problem that we call CliqueHog. Given a graph H, we define a *cliquehog* of size k as a pair (C, F) such that $C \subseteq V(H)$ is a clique of size k, and $F \subseteq E(H)$ is a set of edges that contains exactly two edges connecting c to $V(H) \setminus C$, for each $c \in C$. The input of the CliqueHog problem is a graph H and an integer k, and it asks whether H contains a *cliquehog* of size k.

[8] The input to this problem is a propositional formula φ in conjunctive normal form with only negative literals and clauses of size two, together with an integer parameter k; the question is whether the formula can be satisfied by a variable assignment that sets exactly k variables to true.

[9] The reduction from weak to strict rankings increases the parameter k by a non-constant term. Thus $\mathsf{W}_l[1]$-hardness does not carry over and the parameterized complexity of k-unpopularity margin with strict rankings is still open; see the related open question in Section 8.

Lemma 6.1. *The* CliqueHog *problem is* NP*-hard and* $\mathsf{W}_l[1]$*-hard with parameter* k*.*

Proof. We reduce from the Clique problem. Given a graph $H = (V, E)$ and an integer k, we construct a graph H' by adding $2|V|$ edges and $2|V|$ vertices to H as follows: for each vertex v of H we introduce two new vertices v' and v'', together with the edges (v, v') and (v, v''). It is then easy to see that H contains a clique of size k if and only if H' contains a cliquehog of size k. The reduction is a linear FPT-reduction with parameter k, as well as a polynomial-time reduction. □

Let us now prove Theorem 1.3 by presenting a reduction from the CliqueHog problem.

Construction. Let $H = (V, E)$ and k be our input for CliqueHog. We construct an instance G of the k-unpopularity margin problem, with a set A of agents and a set B of objects as follows. For each $v \in V$, we define a vertex gadget G_v containing agents a_v^0 and a_v^1 and objects b_v^0 and b_v^1. For each $e \in E$, we define an agent a_e and an object b_e. Furthermore, we will use a set A_D of $|E| - \binom{k}{2} - 2k$ dummy agents, and a set B_D of $|E| - \binom{k}{2} - 2k$ artificial objects. We define $A_V = \{a_v^i : v \in V, i \in \{0,1\}\}$, $B_V = \{b_v^i : v \in V, i \in \{0,1\}\}$, $A_E = \{A_e : e \in E\}$, and $B_E = \{b_e : e \in E\}$. We set $A = A_V \cup A_E \cup A_D$ and $B = B_V \cup B_E \cup B_D$. The preferences of the agents in G are as follows (ties are simply indicated by including them as a set in the preference list):

$$
\begin{array}{ll}
a_v^i : \{b_e : e \text{ is incident to } v \text{ in } H\} \succ b_v^0 \succ b_v^1 & \text{for each } v \in V \text{ and } i \in \{0,1\}; \\
a_e : b_e \succ B_D \cup \{b_x^0, b_y^0\} \succ \{b_x^1, b_y^1\} & \text{for each } e = (x,y) \in E; \\
a_d : B_E & \text{for each } a_d \in A_D.
\end{array}
$$

We finish the construction by setting the bound for the unpopularity margin of the desired assignment as k. Clearly, this is a polynomial-time reduction, and also a linear FPT-reduction with parameter k, so it remains to prove that H contains a cliquehog of size k if and only if G admits an assignment M with unpopularity margin at most k.

Lemma 6.2. *If* (C, F) *is a cliquehog in* H *of size* k*, then* G *admits an assignment* M *with unpopularity margin at most* k*.*

Proof. Let f_c^0 and f_c^1 denote the two edges of F connecting c to $V \setminus C$ in H (in any fixed order), and we set $F^i = \{f_c^i : c \in C\}$ for $i \in \{0,1\}$.

We define an assignment M as follows; see Fig. 3 as an illustration. First, let us assign the $|E| - \binom{k}{2} - 2k$ objects in $\{b_e : e \notin E[C] \cup F\}$ to the dummy agents (where $E[C]$ denotes the set of those edges of E whose both endpoints are in C). Second, we assign the artificial objects to the $|E| - \binom{k}{2} - 2k$ agents in $\{a_e : e \notin E[C] \cup F\}$. To define M on the remaining objects and agents, let

$$
\begin{array}{ll}
M(a_v^i) = b_v^i & \text{for each } v \in V \setminus C \text{ and } i \in \{0,1\}; \\
M(a_v^i) = b_{f_v^i} & \text{for each } v \in C \text{ and } i \in \{0,1\}; \\
M(a_e) = b_e & \text{for each } e \in E[C]; \\
M(a_f) = b_v^i & \text{for each } f \in F \text{ where } f = f_v^i.
\end{array}
$$

Observe that M indeed assigns exactly one object to each agent. To show that $\mu(M) \leq k$, we define a dual certificate $\vec{\alpha}$ for M:

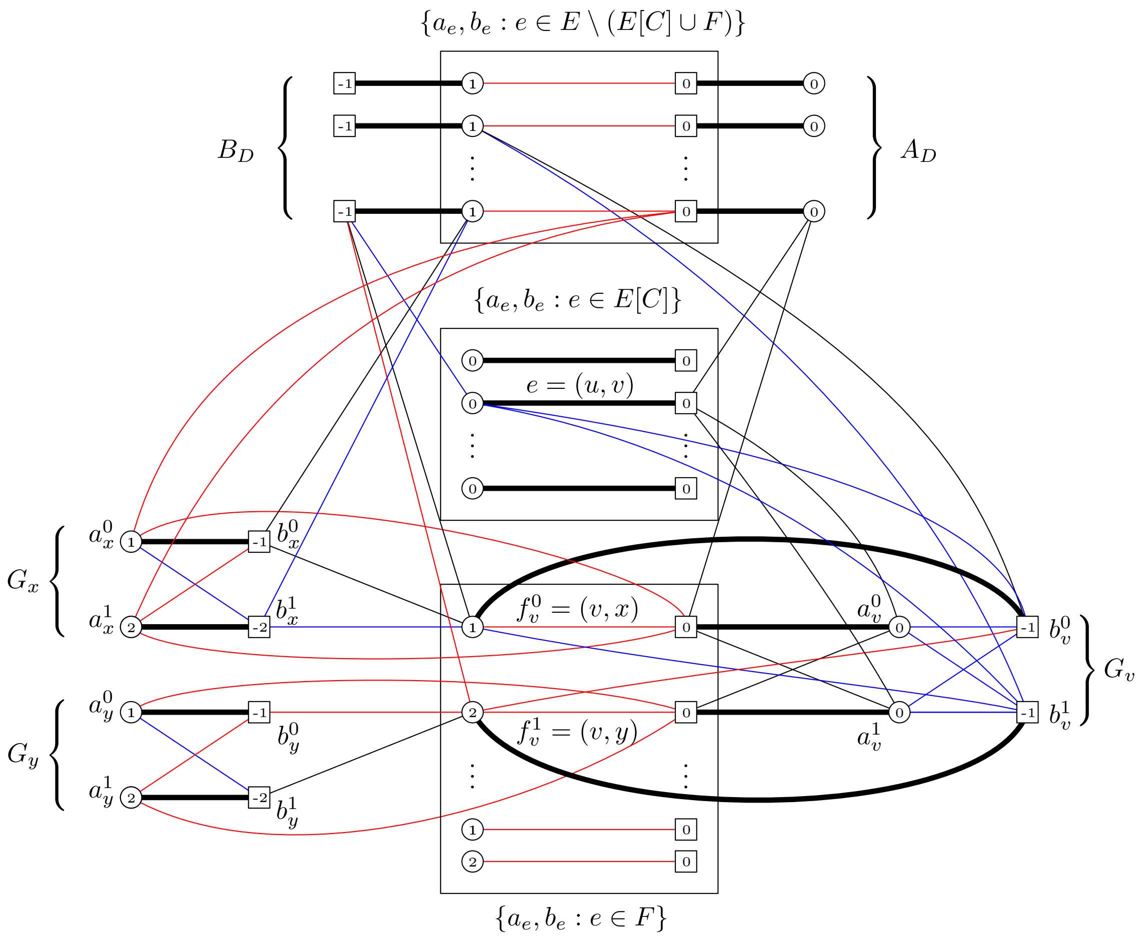

Fig. 3. Illustration for the construction in the proof of Theorem 1.3. The assignment M defined in Lemma 6.2 is indicated by bold lines. Red, black, and blue edges have weight $+1$, 0, and -1, respectively, according to $\mathsf{wt}_M(\cdot)$. The values of the dual certificate $\vec{\alpha}$ for M are indicated by numbers within the circle (square) corresponding to the given agent (object, respectively). For some edges (a_e, b_e) in G we indicate only the corresponding edge e of H (see (u, v), (v, x) and (v, y)). The figure assumes $v \in C$ but $x, y \notin C$; note that the edge $(a_{(v,y)}, b_v^1)$ is overloaded in M.

$\alpha_{a_d} = 0$	for each $a_d \in A_D$;	$\alpha_{b_d} = -1$	for each $b_d \in B_D$;
$\alpha_{a_v^0} = 1$	for each $v \in V \setminus C$;	$\alpha_{b_v^0} = -1$	for each $v \in V \setminus C$;
$\alpha_{a_v^1} = 2$	for each $v \in V \setminus C$;	$\alpha_{b_v^1} = -2$	for each $v \in V \setminus C$;
$\alpha_{a_v^i} = 0$	for each $v \in C$ and $i \in \{0, 1\}$;	$\alpha_{b_v^i} = -1$	for each $v \in C$ and $i \in \{0, 1\}$;
$\alpha_{a_e} = 1$	for each $e \in E \setminus (F^1 \cup E[C])$;	$\alpha_{b_e} = 0$	for each $e \in E$;
$\alpha_{a_e} = 2$	for each $e \in F^1$;		
$\alpha_{a_e} = 0$	for each $e \in E[C]$.		

It is straightforward to check that $\vec{\alpha}$ is indeed a dual certificate, that is, it satisfies the constraints of LP2. Since $\sum_{u \in A \cup B} \alpha_u = k$, assignment M indeed has unpopularity margin at most k. $\square$

We now show that any assignment in G with unpopularity margin at most k implies the existence of a cliquehog of size k in H. We first establish a useful assumption that we will show is without loss of generality. We say that an assignment M has a *nice structure*, if for each $e \in E$ one of the following cases holds true:

- $M(a_e) = b_e$, or
- $M(a_e) = b_x^i$ and $M(a_x^i) = b_e$ for some endpoint x of edge e in H and $i \in \{0, 1\}$, or
- $M(a_e) \in B_D$ and $M(b_e) \in A_D$.

Lemma 6.3. *If there exists an assignment in G with unpopularity margin at most k, then there exists an assignment in G with unpopularity margin at most k that has a nice structure.*

Proof. Let M' be any assignment in G. By switching the names of the identical agents a_x^0 and a_x^1 wherever necessary, we can assume that $M'(a_x^i) \in \{b_x^i\} \cup B_E$ for all $x \in V$ and $i \in \{0, 1\}$. We construct a new assignment M as follows: For each $e \in E$, let $M(a_e) = M'(a_e)$. For each $x \in V$ and each $i \in \{0, 1\}$, if $M'(b_x^i) = a_e$ for some $e \in E$, then let $M(a_x^i) = b_e$; otherwise let $M(a_x^i) = b_x^i$. Assign the unmatched objects in B_E arbitrarily to the dummy agents in A_D. Note that for every agent $a \in A$, we have $M(a) = M'(a)$ or $M(a)$ and $M'(a)$ are both in B_E. Since every agent is indifferent between all her neighbors in B_E, the assignment M has the same unpopularity margin as M'. $\square$

In what follows, let M be an assignment in G with unpopularity margin at most k and a nice structure, and let $\vec{\alpha}$ be a dual certificate for M. We will construct a cliquehog of size k in H.

We define a partition (E_0, E_1, E_2) of those edges e in E for which $M(a_e) \notin B_D$. For any edge $e \in E$, we decide which set of the partition e belongs to using the following procedure.

1. If $M(a_e) = b_e$ and $\alpha_{a_e} + \alpha_{b_e} > 0$, then we put e into E_0.
2. If $M(a_e) = b_e$ and $\alpha_{a_e} + \alpha_{b_e} = 0$, then we put e into E_1.
3. If $M(a_e)$ is an object in G_x or G_y where $e = (x, y)$, then we put e into E_2.

We also define S as the set of those vertices $v \in V$ for which the vertex gadget G_v contains an agent or object that is incident to an overloaded edge of M. Note that since no agent in a vertex gadget is connected to an object in another vertex gadget, and moreover, the overloaded edges of M corresponding to edges in E_0 are not incident to any vertex gadget, we know that

$$|S| + |E_0| \le \mu(M) \le k. \tag{1}$$

We will show that (S, E_2) is a cliquehog of size k. To do so, we first establish two helpful lemmas. The first shows that S contains all nodes $x \in V$ for which at least one of the objects b_x^0 and b_x^1 is matched to an edge agent. This implies that every edge in E_2 is incident to a node in S. The second lemma shows that all endpoints of edges in E_1 are contained in S.

Lemma 6.4. *Let $x \in V$. If $M(b_x^0) \in A_E$ or $M(b_x^1) \in A_E$, then $x \in S$.*

Proof. We distinguish two cases. First, suppose that $M(b_x^0) = a_e$ for some $e \in E$. Recall that this implies $M(a_x^0) = b_e$, because M has a nice structure. From the preferences of agents and the feasibility of $\vec{\alpha}$, we obtain the three inequalities $\alpha_{a_x^0} + \alpha_{b_x^1} \ge -1$ and $\alpha_{M(b_x^1)} + \alpha_{b_x^0} \ge 1$ and $\alpha_{a_e} + \alpha_{b_e} \ge 1$. Adding these inequalities, we observe that at least one of the sums $\alpha_{a_x^0} + \alpha_{b_e}$ or $\alpha_{b_x^1} + \alpha_{M(b_x^1)}$ or $\alpha_{a_e} + \alpha_{b_x^0}$ must be positive, and thus the corresponding edge in M must be overloaded.

Second, suppose that $M(b_x^0) \notin A_E$ but $M(b_x^1) = a_e$ for some $e \in E$. Recall that this implies $M(a_x^0) = b_x^0$ and $M(a_x^1) = b_e$, because M has a nice structure. Again from the preferences of agents and the feasibility of $\vec{\alpha}$, we obtain the three inequalities $\alpha_{a_e} + \alpha_{b_x^0} \geq 1$ and $\alpha_{a_x^0} + \alpha_{b_e} \geq 1$ and $\alpha_{a_x^1} + \alpha_{b_x^1} \geq -1$. Adding these inequalities, we see that at least one of the sums $\alpha_{a_e} + \alpha_{b_x^1}$ or $\alpha_{a_x^0} + \alpha_{b_x^0}$ or $\alpha_{a_x^1} + \alpha_{b_e}$ must be positive, and so the corresponding edge in M must be overloaded. ◻

Lemma 6.5. *If $e = (x, y) \in E_1$, then $\{x, y\} \subseteq S$.*

Proof. We only show that $x \in S$. The proof for $y \in S$ follows by symmetry. If $M(b_x^0) \in A_E$ or $M(b_x^1) \in A_E$, Lemma 6.4 implies that $x \in S$. So we can assume that $M(b_x^0) = a_x^0$ and $M(b_x^1) = a_x^1$ by the nice structure of M. By the preferences of agents and feasibility of $\vec{\alpha}$, we obtain the three inequalities $\alpha_{a_x^0} + \alpha_{b_e} \geq 1$ and $\alpha_{a_x^1} + \alpha_{b_x^0} \geq 1$ and $\alpha_{a_e} + \alpha_{b_x^1} \geq -1$. Adding up these inequalities, we observe that at least one of the sums $\alpha_{a_x^0} + \alpha_{b_x^0}$ or $\alpha_{a_x^1} + \alpha_{b_x^1}$ or $\alpha_{a_e} + \alpha_{b_e}$ must be positive. Because $e \in E_1$, the third expression is equal to 0, and so one of the former two has to be positive. This implies that at least one of the two corresponding edges of M in G_x is overloaded. ◻

We are now ready to prove Lemma 6.6, which together with Lemma 6.2 proves Theorem 1.3.

Lemma 6.6. *If the constructed instance G admits an assignment M with unpopularity margin at most k, then H contains a cliquehog of size k.*

Proof. By Lemma 6.5 we know that any edge $e \in E_1$ must have both of its endpoints in S, yielding

$$|E_1| \leq \binom{|S|}{2}. \tag{2}$$

By Lemma 6.4, each edge $e \in E_2$ must have its agent a_e assigned to an object in a vertex gadget G_v for some $v \in S$. By Inequality (1) there are at most $2|S|$ such objects, so we obtain

$$|E_2| \leq 2|S|. \tag{3}$$

Recall that by construction of G and the definition of the partition (E_0, E_1, E_2), we know

$$\binom{k}{2} + 2k = |E| - |B_D| = |E_1| + |E_2| + |E_0| \leq \binom{|S|}{2} + |S| + k$$

where the inequality follows from combining Inequalities (1), (2), and (3). Hence, by $|S| \leq k$ we obtain $|S| = k$. Moreover, every inequality we used must hold with equality. In particular, this implies $|E_1| = \binom{k}{2}$, which can only happen if there are $\binom{k}{2}$ edges in H (namely, those in E_1) with both of their endpoints in S. In other words, S forms a clique in H. Additionally, (3) must also hold with equality, so $|E_2| = 2|S| = 2k$. Since every edge in E_2 is incident to a vertex of S by Lemma 6.4, but is not contained in $E[S]$ (because $E[S] = E_1 \subseteq E \setminus E_2$), and any $x \in S$ is incident to at most two edges of E_2 (by the definition of E_2), we can conclude that (S, E_2) is a cliquehog of size k. ◻

6.3 A reduction from weak rankings to strict for the k-unpopularity margin problem

Lemma 6.7. *Let G be an instance with weak rankings and n agents. Then we can compute in polynomial time an instance G' with strict rankings and an integer $q \in \mathcal{O}(n)$ such that G admits an assignment with unpopularity margin k if and only if G' admits an assignment with unpopularity margin $k + q$ for any $k \in [n]$.*

Proof. Let G be an instance with weak rankings. We divide the proof into two parts. In the first part we show that we can assume without loss of generality that the set of agents in G can be partitioned into two groups, $A_\succ$ and $A_\sim$, where agents in $A_\succ$ have strict preferences over the objects in their neighborhood and agents in $A_\sim$ are indifferent among all their neighboring objects. Starting from such an instance, we then prove the lemma.

Part I. Starting from a graph G with weak rankings, we create a graph $\hat{G}$ and preferences as follows: For every agent v in G, let $B_v^{(1)} \succ \cdots \succ B_v^{(r_v)}, r_v \in \{1, \ldots, m\}$ be the weak ranking over its neighboring objects. We add agent v to $\hat{G}$, and for each $i \in \{1, \ldots, r_v\}$ we further add a new agent $a_v^{(i)}$ and a new object $b_v^{(i)}$ to $\hat{G}$. The new agent $a_v^{(i)}$ is connected to—and defined to be indifferent among—all objects in $B_v^{(i)} \cup \{b_v^{(i)}\}$. Lastly, we introduce edges from v to all objects in $\bigcup_{i=1}^{r_v} b_v^{(i)}$, with preferences $b_v^{(1)} \succ b_v^{(2)} \succ \cdots \succ b_v^{(r_v)}$. In $\hat{G}$ we can partition the set of nodes into $A_\succ$, containing copies of agents in G, and $A_\sim$, containing the newly introduced agents. Agents in $A_\succ$ have strict preferences and all agents in $A_\sim$ are indifferent among all their neighbors. Below, we show that the two instances are essentially equivalent.

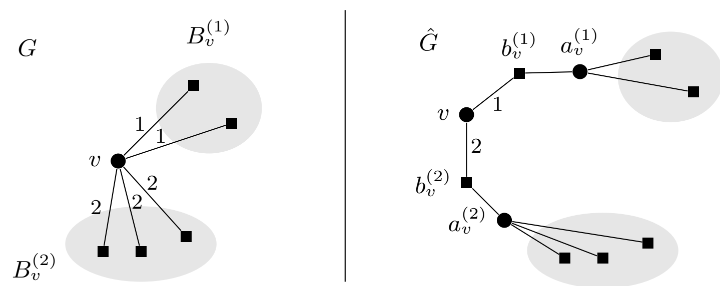

Fig. 4. Illustration of the first part of the proof of Lemma 6.7. The left side illustrates the neighborhood of a fixed node v within the original graph G. The right side illustrates the corresponding situation within the graph $\hat{G}$. Agents are depicted by circles and objects by squares. Labels of the edges indicate the rank of the edge within the ranking of the incident agent. Agents $a_v^{(1)}$ and $a_v^{(2)}$ are indifferent among all their neighbors, hence, their labels are omitted.

More precisely, we show that there exists a bijection f mapping assignments in G to assignments in $\hat{G}$ such that $\Delta(N, M) = \Delta(f(N), f(M))$ for any two assignments M and N in G. Let M be an assignment in G. We start with $f(M) = \emptyset$. For every edge $(v, w) \in M$ we do the following: Let $B_v^{(1)} \succ \cdots \succ B_v^{(r_v)}$ be the preferences of v and let $i \in \{1, \ldots, r_v\}$ be such that $w \in B_v^{(i)}$. Then, we add the edges $(v, b_v^{(i)})$ and $(a_v^{(i)}, w)$ to $f(M)$. Moreover, for all indices $j \in \{1, \ldots, r_v\} \setminus \{i\}$ we add the edge $(a_v^{(j)}, b_v^{(j)})$ to $f(M)$. It is easy to see that this is a bijection.

Since for every agent in $A_\succ$ the rank of its partner in M equals the rank of its partner in $f(M)$ and agents in $A_\sim$ are indifferent among all their neighbors, we get $\Delta(N, M) = \Delta(f(N), f(M))$.

Part II. Due to part I, we can assume that the agents in G are partitioned into sets $A_\succ$ and $A_\sim$ such that agents in $A_\succ$ have strict preferences over objects and the agents in $A_\sim$ are indifferent among all their neighboring objects. Starting from the graph G, we create a graph G' with strict rankings as follows. We first copy all agents and objects to G'. Then for each $a \in A_\sim$ we introduce two new agents a' and a'' and two new objects b'_a and b''_a. We add all possible edges from $\{a, a', a''\}$ to $\mathsf{Nbr}_G(a) \cup \{b'_a, b''_a\}$. The preferences of agents a, a', and a'' are identical, namely, b'_a is their first choice, b''_a their second choice, followed by all objects in $\mathsf{Nbr}_G(a)$ in arbitrary order.

We define a function f which maps assignments in G' to assignments in G. Let M' be an assignment in G'. For each $(a,b) \in M'$ where $a \in A_{\succ}$, we add the edge (a,b) to $f(M')$. Now, observe that for every $a \in A_{\sim}$, exactly one of the nodes a, a', a'' is matched to a node of the original neighborhood, i.e., $\mathsf{Nbr}_G(a)$. Let b be the assigned object. Then we add (a,b) to $f(M')$. It is easy to see that $f(M')$ is then an assignment in G. We continue by observing that

(i) f is surjective, and
(ii) for all $a \in A_{\succ}$ and two assignment M', N' in G': $M' \succ_a N'$ if and only if $f(M') \succ_a f(N')$.

For (i), note that for every assignment M in G we can create an assignment M' in G such that $f(M') = M$ holds, as follows: Copy all edges from M to M' and then add for every $a \in A_{\sim}$ the edges (a', b'_a) and (a'', b''_a). Clearly, $f(M') = M$. For (ii) observe that for all assignments M' in G', agents in $A_{\succ}$ have the same assigned object in M' as in $f(M')$.

We define $q = |A_{\sim}|$ and show that G admits an assignment with unpopularity margin at most k if and only if G' admits an assignment with unpopularity margin at most $k+q$, for any $k \in \mathbb{N}$.

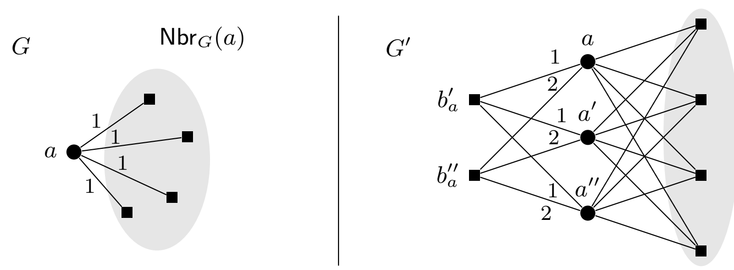

Fig. 5. Illustration of the second part of the proof of Lemma 6.7. The left side illustrates the neighborhood of an agent a, indifferent among all its neighbors, within the graph G. The right side captures the corresponding gadget in the graph G'. Labels on the edges indicate the preferences of agents. The ranks of edges between $\{a, a', a''\}$ and $\mathsf{Nbr}_G(a)$ can be chosen arbitrarily (but need to be larger than 2), hence, these labels are omitted.

Direction "$\Rightarrow$". Assume that G admits an assignment M with unpopularity margin at most k. Choose an assignment M' from G' such that $f(M') = M$ holds (such an M' is guaranteed to exist by (i)). Let N' be an assignment in G' maximizing $\Delta(N', M')$, so $\mu(M') = \Delta(N', M')$. Since $\Delta(f(N'), f(M')) \leq \max_N \Delta(N, M) = k$, we know by (ii) that agents in $A_{\succ}$ contribute at most k to $\Delta(N', M')$. Moreover, we claim that agents not in $A_{\succ}$ contribute at most q to $\Delta(N', M')$. To see this, consider the gadget for some agent $a \in A_{\sim}$. We distinguish two cases. First, assume that the agent from $\{a, a', a''\}$ which is assigned an object in $\mathsf{Nbr}_G(a)$ is the same in assignments M' and N'; w.l.o.g. we assume it is agent a. Then a' and a'' together contribute 0, and a at most 1 to $\Delta(N', M')$. Second, assume w.l.o.g. that a is assigned some object in $\mathsf{Nbr}_G(a)$ by M', while a' is assigned some object in $\mathsf{Nbr}_G(a)$ by N'. Then a and a' together contribute 0,and a'' at most 1. As this holds for every gadget belonging to agents in $A_{\sim}$, this proves $\Delta(N', M') \leq k + q$.

Direction "$\Leftarrow$". For the other direction, assume that G' admits an assignment M' with unpopularity margin at most $k+q$. We claim that $f(M')$ has unpopularity at most k. Assume for contradiction that there exists some assignment N in G with $\Delta(N, f(M')) > k$. We construct N' as follows. For every agent $a \in A_{\succ}$, we let $N'(a) = N(a)$. Next, for every agent $a \in A_{\sim}$, let b_M be its assigned object in $f(M')$ and b_N be its assigned object in N. We can assume w.l.o.g. that a is matched to b_M in M', a' is matched to b'_a, and a'' is matched to b''_a. In N' we match a' to b_N, a'' to b'_a, and a to b''_a. Then a, a', and a'' together contribute exactly 1 to $\Delta(N', M')$. Using the same argument for all $a \in A_\sim$ and (ii) yields that $\mu(M') \geq \Delta(N', M') > k + q$, a contradiction. □

7 The Minimum-Cost Popular Assignment Problem

The minimum-cost popular assignment *problem.* Given a bipartite graph $G = (A \cup B, E)$ where every agent has preferences in partial order over her adjacent objects, together with a cost function $c : E \to \mathbb{R}$ on the edges and a budget $\beta \in \mathbb{R}$, does G admit a popular assignment of total cost at most β?

In this section we prove Theorem 7.1 which shows the minimum-cost popular assignment problem is NP-hard when edge costs are in $\{0, 1, +\infty\}$. This is weaker than Theorem 1.4 which says this problem is NP-hard even when edge costs are in $\{0, 1\}$. The proof of Theorem 7.1 presents the main ideas used in the proof of Theorem 1.4; the latter is a little more involved and can be found in Appendix A.

Theorem 7.1. *The* minimum-cost popular assignment *problem is* NP*-complete, even if each edge cost is in* $\{0, 1, +\infty\}$ *and agents have strict preferences.*

Proof. Since we can check for any assignment of objects to agents, whether it is popular and its cost is within the budget, the problem is clearly in NP. We now present a reduction from the Vertex Cover problem, whose input is a graph $H = (V, E)$ and an integer k, and asks whether there exists a set S of at most k vertices in H such that each edge of E has one of its endpoints in S.

Construction. Let us construct our instance for the minimum-cost popular assignment problem; see Fig. 6. We define the set A of agents and the set B of objects by introducing the following:

- two *level-setting gadgets* G^0_ℓ and G^1_ℓ, with G^i_ℓ consisting of a single edge connecting agent a^i_ℓ and object b^i_ℓ, for each $i \in \{0, 1\}$;
- a *vertex gadget* for each $x \in V$, consisting of a cycle of length 4, containing agents a^0_x and a^1_x, and objects b^0_x and b^1_x;
- an *edge gadget* for each $e \in E$, consisting of a cycle of length 8 containing agents $a^0_e, \dots a^3_e$ and objects $b^0_e, \dots, b^3_e$.

We let A_V and A_E contain agents of all vertex and edge gadgets, respectively, and we define the sets B_V and B_E of objects analogously. Additionally, we introduce a set F of *inter-gadget edges.* First, we add edges of

$$F_\ell = \{(a^1_\ell, b^0_\ell), (a^0_\ell, b^1_\ell)\}$$

to the level-setting gadgets. Second, in order to enforce certain lower bounds on the dual certificate, we connect some agents and objects in the level-setting gadgets with those in the vertex and edge gadgets, by adding the edges of

$$F_{\text{bnd}} = \{(a^0_x, b^0_\ell) : x \in V\} \cup \{(a^0_e, b^1_\ell), (a^3_e, b^1_\ell), (a^1_\ell, b^1_e) : e \in E\}. \tag{4}$$

Third, we encode the incidence relation in H into our instance by adding the edges of

$$F_{\text{inc}} = \{(a^0_e, b^1_y), (a^3_e, b^1_x) : e = (x, y)\} \tag{5}$$

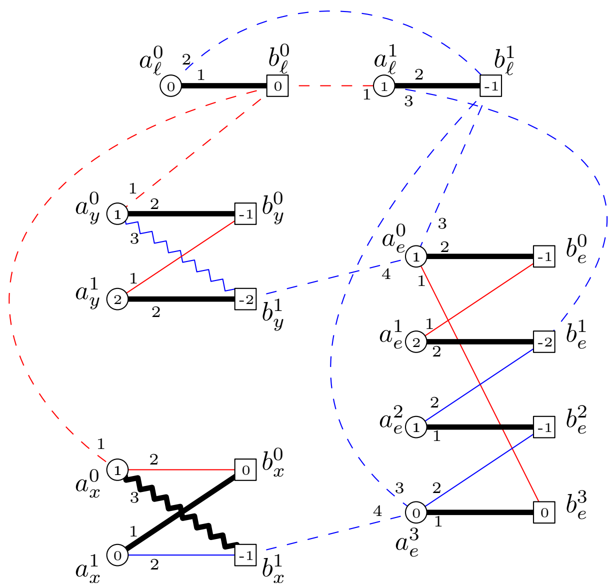

Fig. 6. An illustration of the instance constructed in the proof of Theorem 7.1, showing the level-setting gadgets together with two vertex gadgets corresponding to vertices x and y in H and an edge gadget corresponding to edge $e = (x, y)$. Agents' preferences are indicated by numbers on the edges. Dashed lines represent inter-gadget edges with cost ∞, zigzagged lines represent edges with cost 1, and normal lines represent edges with cost 0. The assignment M defined in direction "$\Leftarrow$" of the proof is indicated by bold lines, assuming a vertex cover S that contains x but not y. Red, black, and blue edges have weight $+1$, 0, and -1, respectively, according to $\mathtt{wt}_M(\cdot)$. The values of the dual certificate $\vec{\alpha}$ for M are indicated by numbers within the circle (square) corresponding to the given agent (object, respectively).

between edge and vertex gadgets. We define the set of all inter-gadget edges as $F = F_\ell \cup F_{\text{bnd}} \cup F_{\text{inc}}$. Note that the edges of F indeed run between different gadgets.

For a set X of objects, we will write $[X]$ in an agent's preference list to denote an arbitrarily ordered list containing objects in X. Then the preferences of the agents are as follows:

$a_\ell^0 : b_\ell^0 \succ b_\ell^1;$
$a_\ell^1 : b_\ell^0 \succ b_\ell^1 \succ \{b_e^1 : e \in E\};$
$a_x^0 : b_\ell^0 \succ b_x^0 \succ b_x^1$ for each $x \in V$;
$a_x^1 : b_x^0 \succ b_x^1$ for each $x \in V$;
$a_e^0 : b_e^3 \succ b_e^0 \succ b_\ell^1 \succ b_y^1$ for each $e = (x, y) \in E$;
$a_e^1 : b_e^0 \succ b_e^1$ for each $e \in E$;
$a_e^2 : b_e^2 \succ b_e^1$ for each $e \in E$;
$a_e^3 : b_e^3 \succ b_e^2 \succ b_\ell^1 \succ b_x^1$ for each $e = (x, y) \in E$;

Finally, we define the cost function: edges of F have cost $+\infty$, the edge (a_x^0, b_x^1) has cost 1 for each $x \in V$, and all remaining edges have cost 0. We set our budget to be k.

We claim that the constructed instance admits a popular assignment of cost at most k if and only if H contains a vertex cover of size at most k.

Direction "$\Rightarrow$". Let M be a popular matching of cost at most k. Since inter-gadget edges have infinite cost, $M(a)$ must be an object in the gadget that contains a, for any agent a. Thus, for any $x \in V$, the cost of the edges of M within the vertex gadget corresponding to x is either 1 (in case M contains the edge (a_x^0, b_x^1)), or 0 (in case M does not contain (a_x^0, b_x^1)). Let S be the set of those vertices $x \in V$ for which the former holds, i.e., $S = \{x \in V : M(a_x^0) = b_x^1\}$; our budget implies $|S| \le k$.

Let M admit a dual certificate $\vec{\alpha}$. Note that w.l.o.g. we can assume that $\alpha_{a_\ell^0} = 0$, as otherwise we can decrease the value of α_a for all agents a by $\alpha_{a_\ell^0}$ and increase α_b by the same amount for each object b. Recall also that for each $(a,b) \in M$ complementary slackness for LP2 implies $\alpha_a + \alpha_b = 0$; hence $\alpha_{b_\ell^0} = -\alpha_{a_\ell^0} = 0$. Since a_ℓ^1 prefers b_ℓ^0 to $M(a_\ell^1) = b_\ell^1$, we know that $\mathsf{wt}_M(a_\ell^1, b_\ell^0) = 1$ and thus $\alpha_{a_\ell^1} \ge 1$ and $\alpha_{b_\ell^1} \le -1$. Using that $\mathsf{wt}_M(a_\ell^0, b_\ell^1) = -1$, we obtain $\alpha_{a_\ell^1} = 1$ and $\alpha_{b_\ell^1} = -1$.

Consider now the edges in F_{bnd}. Observe that $\mathsf{wt}_M(a_e^0, b_\ell^1) = -1$, $\mathsf{wt}_M(a_e^3, b_\ell^1) = -1$ and $\mathsf{wt}_M(a_\ell^1, b_e^1) = -1$ for any $e \in E$. Furthermore, we also have $\mathsf{wt}_M(a_x^0, b_\ell^0) = 1$ for any $x \in V$. These observations imply the following bounds:

$$\min(\alpha_{a_e^0}, \alpha_{a_e^3}) \ge 0 \qquad \text{for each } e \in E; \tag{6}$$

$$\alpha_{b_e^1} \ge -2 \qquad \text{for each } e \in E; \tag{7}$$

$$\alpha_{a_x^0} \ge 1 \qquad \text{for each } x \in V. \tag{8}$$

For some $v \in V \setminus S$, since a_v^1 prefers b_v^0 to $M(a_v^1) = b_v^1$, we know that $\mathsf{wt}_M(a_v^1, b_v^0) = 1$, implying

$$\alpha_{b_v^1} = -\alpha_{a_v^1} \le \alpha_{b_v^0} - 1 = -\alpha_{a_v^0} - 1 \le -2, \tag{9}$$

where the last inequality follows from our bound (8).

Let us fix some edge $e = (x,y) \in E$. Let us define two matchings $M_x^e = \{(a_e^i, b_e^i) : i \in \{0, \dots, 3\}\}$ and $M_y^e = \{(a_e^i, b_e^{(i-1) \bmod 4}) : i \in \{0, \dots, 3\}\}$. Since M does not contain any inter-gadget edges, we know that M contains either M_x^e or M_y^e. Assume first that M contains M_x^e; we claim that $\alpha_{a_e^3} = 0$. For the sake of contradiction, assume otherwise; by (6) this implies $\alpha_{a_e^3} \ge 1$. From the preferences of agents a_e^1 and a_e^0, we know $\mathsf{wt}_M(a_e^1, b_e^0) = 1$ and $\mathsf{wt}_M(a_e^0, b_e^3) = 1$. This implies

$$\alpha_{b_e^1} = -\alpha_{a_e^1} \le \alpha_{b_e^0} - 1 = -\alpha_{a_e^0} - 1 \le \alpha_{b_e^3} - 1 - 1 = -\alpha_{a_e^3} - 2 \le -3, \tag{10}$$

which contradicts our bound (7). Hence, $\alpha_{a_e^3} = 0$. By the preferences of agent a_e^3, we know $\mathsf{wt}_M(a_e^3, b_x^1) = -1$, from which we get $\alpha_{b_x^1} \ge -1$. Since (9) holds for every $v \in V \setminus S$, this implies $x \in S$.

Assume now that M contains M_y^e; we claim that $\alpha_{a_e^0} = 0$. For the sake of contradiction, assume otherwise; by (6) this implies $\alpha_{a_e^0} \ge 1$. From the preferences of agents a_e^2 and a_e^3, we know $\mathsf{wt}_M(a_e^2, b_e^2) = 1$ and $\mathsf{wt}_M(a_e^3, b_e^3) = 1$. This implies

$$\alpha_{b_e^1} = -\alpha_{a_e^2} \le \alpha_{b_e^2} - 1 = -\alpha_{a_e^3} - 1 \le \alpha_{b_e^3} - 1 - 1 = -\alpha_{a_e^0} - 2 \le -3. \tag{11}$$

which contradicts our bound (7). Hence, $\alpha_{a_e^0} = 0$. By the preferences of agent a_e^0, we know $\mathsf{wt}_M(a_e^0, b_y^1) = -1$, from which we get $\alpha_{b_y^1} \ge -1$. Therefore, the bound (9) implies $y \in S$.

Thus, we have proved that $x \in S$ or $y \in S$ holds for any edge $(x,y) \in E$, that is, S is a vertex cover of size at most k in H.

Direction "$\Leftarrow$". For the other direction, given a vertex cover $S \subseteq V$ of size at most k in H, we will show that a popular assignment M of total cost exactly k exists. For each edge $e \in E$, let us fix one of its endpoints in S, and denote it by $\tau(e)$. We may define M as follows:

$$
\begin{array}{ll}
M(a_\ell^i) = b_\ell^i & \text{for any } i \in \{0,1\}, \\
M(a_x^i) = b_x^i & \text{for any } x \in V \setminus S \text{ and } i \in \{0,1\}, \\
M(a_x^i) = b_x^{1-i} & \text{for any } x \in S \text{ and } i \in \{0,1\}, \\
M(a_e^i) = M^e_{\tau(e)}(a_e^i) & \text{for any } e \in E \text{ and } i \in \{0,\dots,3\}.
\end{array}
$$

It is clear that M indeed has total cost k. To show that M is popular, we define a dual certificate for M by defining α_b for each object $b \in B$ as follows; we set $\alpha_a = -\alpha_{M(a)}$ for each agent $a \in A$.

$$
\begin{array}{lllll}
\alpha_{b_\ell^0} = 0; & & \quad & \alpha_{b_\ell^1} = -1; & \\
\alpha_{b_x^0} = 0 & \text{for each } x \in S; & & \alpha_{b_e^0} = -1 & \text{for each } e \in E; \\
\alpha_{b_x^0} = -1 & \text{for each } x \in V \setminus S; & & \alpha_{b_e^1} = -2 & \text{for each } e \in E; \\
\alpha_{b_x^1} = -1 & \text{for each } x \in S; & & \alpha_{b_e^2} = -1 & \text{for each } e \in E; \\
\alpha_{b_x^1} = -2 & \text{for each } x \in V \setminus S; & & \alpha_{b_e^3} = 0 & \text{for each } e \in E.
\end{array}
$$

This finishes the proof of the theorem. □

Minimum-cost Popular Assignment vs. Popular Assignment with Forbidden Edges. The popular assignment with forbidden edges problem can be seen as the special case of minimum-cost popular assignment in which the popular assignment may only contain edges of cost 0, excluding all edges of non-zero cost. In the general version of minimum-cost popular assignment, however, there is a degree of freedom as to which non-zero cost edges are included in the assignment. Our proof of Theorem 7.1 shows that this degree of freedom introduces an additional complexity to the problem. Indeed, in our reduction from Vertex Cover, a set of vertices chosen as a vertex cover is encoded via the set of cost 1 edges chosen to be included in the allocation, and explicitly forbidding all cost 1 edges would turn the constructed popular assignment with forbidden edges instance into a trivial 'no'-instance.

8 Open Problems

We proposed a polynomial-time algorithm for computing a popular assignment in an instance $G = (A \cup B, E)$ with one-sided preferences, if one exists. The running time of our algorithm is $O(m \cdot n^{5/2})$ where $|A| = |B| = n$ and $|E| = m$. Our algorithm solves $O(n^2)$ instances of the maximum matching problem in certain subgraphs of G. It is easy to show instances where our algorithm indeed solves $\Theta(n^2)$ instances of the maximum matching problem. Can we do this more efficiently? Is there a faster algorithm for the popular assignment problem?

Another open problem is to show a short *witness* that a given instance G does not admit a popular assignment. Rather than run our algorithm and discover that G has no popular assignment, is there a forbidden structure that causes G not to admit a popular assignment? Can we characterize instances that admit popular assignments? Interestingly, such a result is known for the stable roommates problem [27]; recall our discussion in Section 1 on the similarity between results for the popular assignment problem and the stable roommates problem.

Our $\mathsf{W}_l[1]$-hardness proof for the k-unpopularity margin problem with parameter k needs weak rankings. We are able to show that this problem remains NP-hard for strict rankings (see Lemma 6.7 in Section 6.3). However, the parameterized complexity of this case is still open: is the k-unpopularity margin problem in FPT with parameter k, if agents' preferences are strict rankings?

Acknowledgments. Telikepalli Kavitha is supported by the DAE, Government of India, under project no. RTI4001. Tamás Király is supported by the National Research, Development and Innovation Fund of Hungary, grant no. K120254 and Application Domain Specific Highly Reliable IT Solutions project under Thematic Excellence Programme TKP2020-NKA-06. Jannik Matuschke is supported by internal funds of KU Leuven. Ildikó Schlotter was supported by the Hungarian Academy of Sciences under its Momentum Programme (LP2016-3/2018), and by the Hungarian Scientific Research Fund through NFKIH grants no. K128611 and K124171. Ulrike Schmidt-Kraepelin is supported by the Deutsche Forschungsgemeinschaft (DFG) under grant BR 4744/2-1.

References

1. A. Abdulkadiroğlu and T. Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66: 689-701, 1998.
2. D. J. Abraham, R. W. Irving, T. Kavitha, and K. Mehlhorn. Popular matchings. *SIAM Journal on Computing*, 37(4): 1030–1045, 2007. (Conference version in SODA 2005.)
3. N. Benabbou, M. Chakraborty, X.-V. Ho, J. Sliwinski, and Y. Zick. Diversity constraints in public housing allocation. In *Proc. of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2018.
4. P. Biró, D. F. Manlove, and S. Mittal. Size versus stability in the marriage problem. *Theoretical Computer Science*, 411: 1828–1841, 2010.
5. J. Chen, X. Huang, I. A. Kanj, and G. Xia. Strong computational lower bounds via parameterized complexity. *Journal of Computer and System Sciences*, 72(8): 1346–1367, 2006.
6. M. Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. L'Imprimerie Royale, 1785.
7. Á. Cseh. Popular matchings. In Ulle Endriss (ed.), *Trends in Computational Social Choice*, 105-121, 2017.
8. Á. Cseh, C.-C. Huang, and T. Kavitha. Popular Matchings with Two-Sided Preferences and One-Sided Ties. *SIAM Journal on Discrete Mathematics*, 31(4): 2348–2377, 2017.
9. R. G. Downey and M. R. Fellows. *Parameterized Complexity.* Monographs in Computer Science, Springer Verlag, 1999.
10. ET Bureau. Final year medics, nurses to be roped in to fight Covid. *Economic Times*, India, May 03, 2021.
11. T. Feder. A new fixed point approach for stable networks and stable marriages. *Journal of Computer and System Sciences*, 45:233–284, 1992.
12. T. Fleiner, R. W. Irving, and D. F. Manlove. Efficient algorithms for generalized stable marriage and roommates problems. *Theoretical Computer Science*, 381(1-3): 162–176, 2007.
13. D. Gale and L. S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69: 9–15, 1962.
14. P. Gärdenfors. Match making: Assignments based on bilateral preferences. *Behavioural Science*, 20:166–173, 1975.
15. R. W. Irving. An efficient algorithm for the stable roommates problem. *Journal of Algorithms*, 6: 577–595, 1985.
16. T. Kavitha. A size-popularity tradeoff in the stable marriage problem. *SIAM Journal on Computing*, 43(1): 52–71, 2014.
17. T. Kavitha. Popular half-integral matchings. In *Proc. of the 43rd International Colloquium on Automata, Languages, and Programming (ICALP)*, 22:1–22:13, 2016.
18. T. Kavitha. Maximum matchings and popularity. In *Proc. of the 48th International Colloquium on Automata, Languages, and Programming (ICALP)*, 85:1–85:21, 2021.
19. T. Kavitha, J. Mestre, and M. Nasre. Popular mixed matchings. *Theoretical Computer Science*, 412: 2679–2690, 2011.
20. T. Kavitha and M. Nasre. Optimal popular matchings. *Discrete Applied Mathematics*, 157(14): 3181–3186, 2009.
21. M. McCutchen. The least-unpopularity-factor and least-unpopularity-margin criteria for matching problems with one-sided preferences. In *Proc. of the 8th Latin American Symposium on Theoretical Informatics (LATIN)*, 593–604, 2016.
22. E. McDermid and R. W. Irving. Popular matchings: Structure and algorithm. *Journal of Combinatorial Optimization*, 22: 339–358, 2011.
23. J. Mestre. Weighted popular matchings. *ACM Transactions on Algorithms*, 10(1): 2, 2014.

24. S. Merrill and B. Grofman. *A Unified Theory of Voting: Directional and Proximity Spatial Models.* Cambridge University Press, 1999.
25. U. G. Rothblum. Characterization of stable matchings as extreme points of a polytope. *Mathematical Programming*, 54: 57–67, 1992.
26. C. T. S. Sng and D. F. Manlove. Popular matchings in the weighted capacitated house allocation problem. *Journal of Discrete Algorithms*, 8(2): 102–116, 2010.
27. J. J. M. Tan. A necessary and sufficient condition for the existence of a complete stable matching. *Journal of Algorithms*, 12: 154–178, 1991.

A Appendix: Proof of Theorem 1.4

We will modify the reduction presented in the proof of Theorem 7.1. Given the instance G constructed there, we obtain a modified instance G' as follows; see Fig. 7 for an illustration.

Construction. We keep each vertex and edge gadget of G, but instead of the two level-setting gadgets in G, we define a single level-setting gadget in G', consisting of a cycle of length $4k'$ where k' is the smallest integer with $2k'-1 > k$. This gadget will contain agents a_ℓ^i for $i \in \{0, \dots, 2k'-1\}$ together with objects b_ℓ^i for $i \in \{0, \dots, 2k'-1\}$. We define the set F of inter-gadget edges by keeping those inter-gadget edges defined in G that run between different gadgets in G'. Hence, we let $F = F_{\text{bnd}} \cup F_{\text{inc}}$ where F_{bnd} and F_{inc} are defined as in (4) and (5), respectively.

The preferences of any agent in a vertex or edge gadget of G' are the same as in G. We define the preferences of the agents in the level-setting gadget as follows:

$a_\ell^0 : b_\ell^0 \succ b_\ell^{2k'-1}$;
$a_\ell^1 : b_\ell^0 \succ b_\ell^1 \succ \{b_e^1 : e \in E\}$;
$a_\ell^i : b_\ell^{i-1} \succ b_\ell^i$ for each $i \in \{2, \dots, k'\}$;
$a_\ell^i : b_\ell^i \succ b_\ell^{i-1}$ for each $i \in \{k'+1, \dots, 2k'-1\}$;

We set the cost of edge (a_ℓ^i, b_ℓ^{i-1}) for each $i \in \{0, \dots, 2k'-1\}$ as 1; we interpret superscripts of objects belonging to the level-setting gadget modulo $2k'$ (here and later). We also set the cost of the edge (a_x^0, b_x^1) for each $x \in V$ as 1. All remaining edges have cost 0, and we set k as our budget.

Direction "⇒". To prove the correctness of our reduction, first we will show that the existence of a popular assignment M of cost at most k in G' implies a vertex cover of size at most k in H.

Excluding inter-gadget edges. We claim that M does not contain any inter-gadget edges. To prove this, we first show that $(a_\ell^0, b_\ell^0) \in M$. Observe that for any $i \in \{0\} \cup \{2, \dots, 2k'-1\}$, agent a_ℓ^i can be assigned either b_ℓ^i or b_ℓ^{i-1}. Using that M assigns an object to every agent, we immediately get that $(a_\ell^0, b_\ell^0) \notin M$ implies that M must contain all of the edges (a_ℓ^i, b_ℓ^{i-1}) for $i \in \{0\} \cup \{2, \dots, 2k'-1\}$. However, these edges have total cost $2k'-1 > k$ which exceeds our budget. Hence, M contains (a_ℓ^0, b_ℓ^0) as well as every edge (a_ℓ^i, b_ℓ^i) for $i \in \{2, \dots, 2k'-1\}$.

Consider now the two agents in some vertex gadget: since neither of them can obtain object b_ℓ^0, they must be assigned the two objects within their gadget.

It remains to show that $(a_\ell^1, b_\ell^1) \in M$. The only other possibility is that $M(a_\ell^1) = b_e^1$ for some edge $e \in E$. This immediately implies $M(a_e^1) = b_e^0$ and $M(a_e^2) = b_e^2$. Let $e = (x, y)$, and let us define $M_x^e = \{(a_e^i, b_e^i) : i \in \{0, \dots, 3\}\}$ and $M_y^e = \{(a_e^i, b_e^{i-i}) : i \in \{0, \dots, 3\}\}$. Note that the objects available for the agents a_e^0 and a_e^3 are b_e^3 and b_ℓ^1 (since all objects of a vertex gadget are assigned within their gadget). Therefore, we have two cases:

a) $(a_e^0, b_\ell^1) \in M$, implying $M(a_e^3) = b_e^3$. Then the assignment which uses the edges of M_x^e and the edge (a_ℓ^1, b_ℓ^1) and otherwise coincides with M is more popular than M, a contradiction.

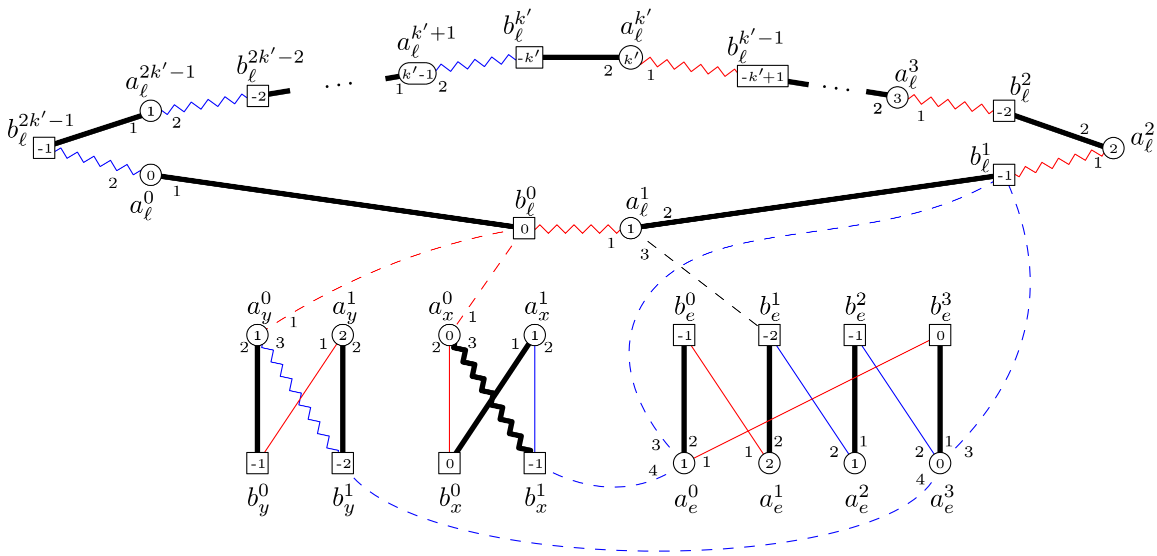

Fig. 7. An illustration of the instance constructed in the proof of Theorem 1.4. As before, agents' preferences are indicated by numbers on the edges. Zigzagged lines represent edges with cost 1, all other edges have cost 0. Dashed lines represent inter-gadget edges. The assignment M defined in direction "$\Leftarrow$" of the proof is indicated by bold lines, assuming a vertex cover S that contains x but not y. Red, black, and blue edges have weight $+1$, 0, and -1, respectively, according to $\mathsf{wt}_M(\cdot)$. The values of the dual certificate $\vec{\alpha}$ for M are indicated by numbers within the circle (square) corresponding to the given agent (object, respectively).

b) $(a_e^3, b_\ell^1) \in M$, implying $M(a_e^0) = b_e^3$. Then the assignment which uses the edges of M_y^e and the edge (a_ℓ^1, b_ℓ^1) and otherwise coincides with M is more popular than M, a contradiction.

Thus, we have proved $(a_\ell^1, b_\ell^1) \in M$, showing that M indeed avoids all inter-gadget edges.

Defining a vertex cover. We are going to define a set S the same way as we did for the instance G, and show that S is a vertex cover in H. Hence, consider any $x \in V$. The cost of the edges of M within the vertex gadget corresponding to x is either 1 (in case M contains the edge (a_x^0, b_x^1)), or 0 (in case M does not contain (a_x^0, b_x^1)). Let S be the set of those vertices $x \in V$ for which the former holds, i.e., $S = \{x \in V : M(a_x^0) = b_x^1\}$; our budget implies $|S| \le k$.

Let M admit a dual certificate $\vec{\alpha}$. Note that w.l.o.g. we can assume that $\alpha_{a_\ell^0} = 0$, as otherwise we can decrease the value of α_a for all agents a by $\alpha_{a_\ell^0}$ and increase α_b by the same amount for each object b. Recall also that $M(a_\ell^i) = b_\ell^i$ for each $i \in \{0, \ldots, 2k'-1\}$. For each $(a, b) \in M$ complementary slackness for LP2 implies $\alpha_a + \alpha_b = 0$, so $\alpha_{b_\ell^i} = -\alpha_{a_\ell^i}s$.

Since a_ℓ^i prefers b_ℓ^{i-1} to $M(a_\ell^i) = b_\ell^i$ for any $i \in \{1, \ldots, k'\}$, we know $\mathsf{wt}_M(a_\ell^i, b_\ell^{i-1}) = 1$. Using this iteratively for $i = 1, 2, \ldots, k'$ we get that $\alpha_{a_\ell^i} \ge i$ and $\alpha_{b_\ell^i} \le -i$ for any $i \in \{1, \ldots, k'\}$. Similarly, using iteratively that $\mathsf{wt}_M(a_\ell^{2k'-i+1}, b_\ell^{2k'-i}) = -1$ for $i = 1, 2, \ldots, k'$, we obtain that $\alpha_{b_\ell^{2k'-i}} \ge -i$ and $\alpha_{a_\ell^{2k'-i}} \le i$. Now, considering the above two observations regarding $\alpha_{a_\ell^{k'}}$ we can conclude that only $\alpha_{a_\ell^{k'}} = k'$ is possible. Moreover, this implies that each of the above inequalities must hold with equality, that is,

$$\alpha_{a_\ell^i} = \begin{cases} i & \text{if } 0 \le i \le k', \\ 2k' - i & \text{if } k' \le i \le 2k' - 1. \end{cases} \tag{12}$$

Consider now the edges in F_{bnd}. Observe that $\mathtt{wt}_M(a_e^0, b_\ell^1) = -1$, $\mathtt{wt}_M(a_e^3, b_\ell^1) = -1$ and $\mathtt{wt}_M(a_\ell^1, b_e^1) = -1$ for any $e \in E$. Furthermore, we also have $\mathtt{wt}_M(a_x^0, b_\ell^0) = 1$ for any $x \in V$. Taking into account (12), these observations yield the following bounds:

$$\min(\alpha_{a_e^0}, \alpha_{a_e^3}) \ge 0 \qquad \text{for each } e \in E; \tag{13}$$

$$\alpha_{b_e^1} \ge -2 \qquad \text{for each } e \in E; \tag{14}$$

$$\alpha_{a_x^0} \ge 1 \qquad \text{for each } x \in V. \tag{15}$$

Notice that these are exactly the same bounds we obtained for the instance I in the proof of Theorem 7.1 in Inequalities (6), (7), and (8). Therefore, using the same arguments again, we obtain that S is a vertex cover of size at most k.

Direction "$\Leftarrow$". For the other direction, given a vertex cover $S \subseteq V$ of size at most k in H, we will show that a popular assignment M of total cost exactly k exists in G'. For each edge $e \in E$, let us fix one of its endpoints in S, and denote it by $\tau(e)$. We may define M as follows:

$$\begin{array}{ll} M(a_\ell^i) = b_\ell^i & \text{for any } i \in \{0, 1, \ldots, 2k' - 1\}, \\ M(a_x^i) = b_x^i & \text{for any } x \in V \setminus S \text{ and } i \in \{0, 1\}, \\ M(a_x^i) = b_x^{1-i} & \text{for any } x \in S \text{ and } i \in \{0, 1\}, \\ M(a_e^i) = M_{\tau(e)}^e(a_e^i) & \text{for any } e \in E \text{ and } i \in \{0, \ldots, 3\}. \end{array}$$

It is clear that M indeed has total cost k. To show that M is popular, we define a dual certificate for M by defining α_b for each object $b \in B$ as follows; we set $\alpha_a = -\alpha_{M(a)}$ for each agent $a \in A$.

$$\begin{array}{llll} \alpha_{b_\ell^i} = -i & \text{for each } i \in \{0, 1, \ldots, k'\}; & \alpha_{b_\ell^{2k'-i}} = -i & \text{for each } i \in \{1, \ldots, k' - 1\}; \\ \alpha_{b_x^0} = 0 & \text{for each } x \in S; & \alpha_{b_e^0} = -1 & \text{for each } e \in E; \\ \alpha_{b_x^0} = -1 & \text{for each } x \in V \setminus S; & \alpha_{b_e^1} = -2 & \text{for each } e \in E; \\ \alpha_{b_x^1} = -1 & \text{for each } x \in S; & \alpha_{b_e^2} = -1 & \text{for each } e \in E; \\ \alpha_{b_x^1} = -2 & \text{for each } x \in V \setminus S; & \alpha_{b_e^3} = 0 & \text{for each } e \in E. \end{array}$$

This finishes the proof of the theorem. $\square$