ECAI 2023
K. Gal et al. (Eds.)

doi:10.3233/FAIA230262

On the Manipulability of Maximum Vertex-Weighted Bipartite b-Matching Mechanisms

Gennaro Auricchio[a;*] **and Jie Zhang**[a]

[a]Department of Computer Science, University of Bath

Abstract. In this paper, we study the Maximum Vertex-weighted b-Matching (MVbM) problem on bipartite graphs in a new game-theoretical environment. In contrast to other game-theoretical settings, we consider the case in which the value of the tasks is public and common to every agent so that the private information of every agent consists of edges connecting them to the set of tasks. In this framework, we study three mechanisms. Two of these mechanisms, namely $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$, are optimal but not truthful, while the third one, $\mathbb{M}_{AP}$, is truthful but sub-optimal. Albeit these mechanisms are induced by known algorithms, we show $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ are the best possible mechanisms in terms of Price of Anarchy and Price of Stability, while $\mathbb{M}_{AP}$ is the best truthful mechanism in terms of approximated ratio. Furthermore, we characterize the Nash Equilibria of $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ and retrieve sets of conditions under which $\mathbb{M}_{BFS}$ acts as a truthful mechanism, which highlights the differences between $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$. Finally, we extend our results to the case in which agents' capacity is part of their private information.

1 Introduction

Managers of large companies are periodically required to submit a list of the projects they carried out for external evaluations. Alongside the list of projects, a manager needs to identify a worker liable for the achievement of every project. Due to the workload cap, every worker can be found responsible for only a finite number of projects. Moreover, every project has a different prestige, which can be regarded as its overall value. The manager aims to maximize the total value of the reported projects by deploying its staff following the aforementioned constraints. Meanwhile, every worker is interested in being associated with high-value projects rather than low-value ones. Therefore, there might be workers that benefit by hiding part of their connections to low value projects. A similar situation occurs in the universities of several European countries. Indeed, to assess the research impact of their higher education institutions, the country asks the universities for a list of their best publications, along with the name of the main author.[1] On one hand, every university wants to find a matching between its lecturers and the publications that maximize the total impact. On the other hand, individual academics want to be designated as the main author of their best publications. In both scenarios, we need to allocate a set of resources that does possess an objective and publicly known value, be they projects and their prestige or publications and their impact score, to a set of self-interested agents.

Aside from these two examples, there are many other real-life situations that can be rephrased as matching problems with self-interested agents. Matching problems were first introduced to minimize transportation costs [30, 34] and to optimally assign workers to job positions [20, 43]. Thereafter, bipartite matching found application in several applied problems, such as sponsored searches [11, 41], school admissions [1, 12], scheduling [46, 29], and general resource allocation [24, 44]. Indeed, characterizing matching through the maximization or minimization of a functional allows to define mathematical objects that arise in several applied contexts. For example, Maximum Cardinality Matchings have connections with the computation of perfect matchings [26], bottleneck matchings [21], and Lévy-Prokhorov distances, [37]. Another example is the Maximum Edge-weighted Matching problem, which has been widely used in clustering problems [18], machine learning [10], and also to compute Wasserstein-barycenters [5]. For a complete discussion of the matching problems and their applications, we refer to [38].

Game-theoretic aspects of matching problems. Matching problems have been extensively studied from a game-theoretical perspective in various contexts, including one-sided matching (such as house allocation [31, 17] and the resident/hospital problem [39]) and two-sided matching (such as stable marriage [27], student admission [9], PhD grants assignment [14], and market matching [35]). Numerous variants of these basic models have been proposed, often with meaningful constraints such as regional constraints [32, 28], diversity constraints [40, 8, 16, 22], or lower quotas [13, 45]. Perhaps the most generic variation on this framework has been studied in [23] and [6], where the authors adapt the game-theoretical framework to the class of Generalized Assignment Problems. In every setting containing self-interested agents, it is important to assess the impact that the mechanism has on the social problem. Consequently, several papers have examined the social aspects of these mechanisms, including manipulability [25, 36, 15], fairness [33], and envy-freeness [3, 2]. For a comprehensive survey of the game-theoretical properties of mechanisms in matching problems, we refer the reader to [7]. All these works, however, assume that the values of the tasks are subjective, meaning that the same task may have a different value for different agents. Yet, in many cases, the value of the objects to be assigned is publicly known. Thus, we cannot model the agents' private information as their ordinal or cardinal preferences over the set of tasks. In this paper, we focus on this distinct class of problems.

Our Contributions. In this paper, we consider a game-theoretical framework in which the agents' private information is the set of edges

* Corresponding Author. Email: ga647@bath.ac.uk
[1] This, for example, is a common practice in the United Kingdom, see https://www.ref.ac.uk for a reference.

Table 1. The mechanisms we study and their properties in the two game-theoretical settings we consider . The "Yes*" indicates that the mechanism is group strategyproof under minor assumptions.

	Truthful	Group SP	Optimal	Efficiency
$\mathbb{M}_{BFS}$	No	No	Yes	PoA = PoS = 2
$\mathbb{M}_{DFS}$	No	No	Yes	PoA = PoS = 2
$\mathbb{M}_{AP}$	Yes	Yes*	No	ar = 2

connecting the agent to the tasks. We assume that the values of the tasks are public and common to every agent. Moreover, we assume that the agents are bounded by their statements so that agents can hide edges, but they cannot report non-existing edges. In this framework, we study three mechanisms induced by known algorithms. Two of them, namely $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$, are induced by the algorithm proposed in [42]. The third one, namely $\mathbb{M}_{AP}$, is induced by the algorithm proposed in [19], which is an approximation of the one proposed in [42]. First, we study the truthfulness and the group strategyproofness of these mechanisms. Both $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ are optimal, but not truthful nor group strategyproof. On the contrary, $\mathbb{M}_{AP}$ is not optimal, but truthful and, under mild assumptions, also group strageyproof. We show that $\mathbb{M}_{AP}$ has an approximation ratio (ar) equal to 2 and that this value is tight, i.e. no other deterministic and truthful mechanism can beat this constant value. We then study the Nash Equilibrium induced by $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$, we use these characterizations to prove that the Price of Anarchy (PoA) and the Price of Stability (PoS) of both mechanisms is equal to 2. As for $\mathbb{M}_{AP}$, these efficiency guarantees are tight: there does not exist a mechanism that can achieve a smaller PoA or PoS. We then characterize some classes of inputs on which $\mathbb{M}_{BFS}$ is truthful, while $\mathbb{M}_{DFS}$ is not. In particular, we infer that, despite their similarities, $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ do possess different game-theoretical properties. Finally, we extend our results to the case in which the agents are able to report their capacity along with their edges. In Table 1, we summarize our main findings.

Due to space limitations, we defer part of the proofs to the Appendix.[2]

2 Preliminaries

In this section, we recall the basic notions on the MVbM problem and introduce the algorithm defining the mechanisms we study.

The Maximum Vertex-weighted b-Matching Problem. Let $G = (A \cup T, E)$ be a bipartite graph. Throughout the paper we refer to $A = \{a_1, a_2, \cdots, a_n\}$ as the set of agents and refer to $T = \{t_1, t_2, \cdots, t_m\}$ as the set of tasks. The set E contains the edges of the bipartite graph. We say that an edge $e \in E$ belongs to agent a_i if $e = (a_i, t_j)$ for a $t_j \in T$. We denote with $M = |E|$ and $N = |A \cup T| = n + m$ the number of edges and the total number of vertices of the graph, respectively. Since the graphs are undirected, we denote an edge by (a_i, t_j) and (t_j, a_i) interchangeably. Moreover, for any given agent $a_i \in A$, we define the set $T_{E,i}$ as the set of tasks in T that are connected to agent a_i through the set of edges E. When it is clear from the context which set E we are referring to, we drop the subscript from $T_{E,i}$ and use T_i. Let $\boldsymbol{b} = (b_1, b_2, \dots, b_n)$ be the vector containing the capacities of the agents, where $b_i \in \mathbb{N}$ for every $i = 1, \dots, n$. A subset $\mu \subset E$ is a b-matching if, for every vertex $a_i \in A$, the number of edges in μ linked to a_i is less than or equal to b_i and, for every vertex $t_j \in T$, the number of edges in μ linked to t_j is, at most, one. Given a b-matching μ, the vertex $a_i \in A$ is *saturated* with respect to μ if the number of edges in μ linked to a_i is exactly b_i. Otherwise, the vertex a_i is *unsaturated* with respect to μ. We denote with $\boldsymbol{q} = (q_1, q_2, \dots, q_m)$ the vector containing the values of the tasks, in our framework, we have $q_j > 0$ for every j. The value of a matching μ is $w(\mu) := \sum_{t_j \in T_\mu} q_j$, where $T_\mu \subset T$ is the set of tasks matched by μ. Given a bipartite graph and a vector $\boldsymbol{b}$, the MVbM problem consists in finding a b-matching μ that maximizes $w(\mu)$.

Finally, given two sets of edges μ_1 and μ_2, denote with $\mu_1 \oplus \mu_2 = (\mu_1 \backslash \mu_2) \cup (\mu_2 \backslash \mu_1)$ their *symmetric difference*. A path $P = \{(t_{j_1}, a_{i_1}), (a_{i_1}, t_{j_2}), (t_{j_2}, a_{i_2}), \cdots, (t_{j_L}, a_{i_L})\}$ in G is a sequence of edges that joins a sequence of vertices. We say that P has a length equal to λ if it contains λ edges. Given a path P and a b-matching μ, P is an *augmenting path* with respect to μ if every vertex a_{i_ℓ}, for $\ell = 1, \dots, L-1$, is saturated with respect to μ, $a_{i_L} \in A$ is unsaturated with respect to μ, and the edges in the path alternatively do not belong to μ and belong to μ. That is, $(t_{j_\ell}, a_{i_\ell}) \notin \mu$, $\ell \in [L]$ and $(a_{i_\ell}, t_{j_{\ell+1}}) \in \mu$, $\ell \in [L-1]$, where $[L]$ is the set containing the first L natural numbers, i.e. $[L] := \{1, 2, \dots, L\}$.

Algorithm 1 Algorithm for MVbM, [42]

Input: A bipartite graph $G = (A \cup T, E)$; agents' capacities $b_i, i = 1, \cdots, n$; task weights $q_j, j = 1, \cdots, m$.
Output: An MVbM μ.

```
μ_0 ← ∅;
Sort q_j, j = 1, ..., m, in decreasing order;
for each j ∈ T do
    if there is augmenting path P starting from j w.r.t. μ_{j−1} then
        μ_j = μ_{j−1} ⊕ P;
    else
        μ_j = μ_{j−1};
    end if
end for
return μ = μ_m;
```

The Algorithm Outline. In this paper, we consider two well-known algorithms from a mechanism design perspective. The first algorithm is presented in [42] and its routine consists in defining a sequence of matchings, namely $\{\mu_j\}_{j=0,1,\dots,m}$, of increasing cardinality. Indeed, the first matching is set as $\mu_0 = \emptyset$ and, given μ_{j-1}, μ_j is defined as follows: if there exists P_j an augmenting path (with respect to μ_{j-1}) that starts from t_j, then $\mu_j = \mu_{j-1} \oplus P_j$, otherwise $\mu_j = \mu_{j-1}$. In Algorithm 1, we sketch the routine of the algorithm.

In [42] it has been shown that Algorithm 1 returns a solution to the MVbM problem in $O(NM)$ time, regardless of whether we find the augmenting path using the *Breadth-First Search* (BFS) or the *Depth-First Search* (DFS). Both the DFS and the BFS when they traverse the graph in search of an augmenting path implicitly assume that there is an ordering of the agents. We say that agent a_i has a higher *priority* than agent a_j if a_i is explored before a_j. In particular, a_1 has the highest priority. We say that Algorithm 1 is endowed with the BFS (or endowed with the DFS) if we use the BFS (or DFS) to find an augmenting path.

The second algorithm is introduced in [19] and [4] and is an approximation version of Algorithm 1 that searches only among the augmenting paths whose length is 1. The authors showed that this approximation algorithm finds a matching whose weight is, in the worst case, half the weight of the MVbM. Notice that for this approximation algorithm, using the BFS or the DFS does not change the outcome thus we omit which traversing graph method is used.

[2] See full version at https://arxiv.org/abs/2307.12305.

3 The Game-Theoretical framework

In this section, we formally define the space of the agents' strategies and the mechanisms we study throughout the paper.

The Strategy Space of the Agents. Given a bipartite graph $G = (A \cup T, E)$, a capacity vector $\boldsymbol{b}$, a value vector $\boldsymbol{q}$, and a b-matching μ over G, we define the social welfare achieved by μ as the total value of the tasks assigned to the agents. Since the social welfare achieved by μ is equal to the total weight of the matching, we use $w(\mu)$ to denote it. We then define the utility of agent a_i as $w_i(\mu) := \sum_{t_j \in T_{\mu,i}} q_j$, where $T_{\mu,i}$ is the set of tasks that μ matches with agent a_i. Notice that $w(\mu) = \sum_{i=1}^{n} w_i(\mu)$.

In this paper, we focus on two settings: (i) The private information of each agent consists of the set of edges that connect the agent to the tasks' set. We call this setting Edge Manipulation Setting (EMS). (ii) The private information of each agent consists of the set of edges and its own capacity. We call this setting Edge and Capacity Manipulation Setting (ECMS). In both cases, we assume that every agent is bounded by its statement, thus every agent is able to report incomplete information, but they are not allowed to report false information. In EMS, this means that an agent can hide some of the edges that connect it to the tasks, but it cannot report an edge that does not exist. The set of strategies of agent a_i, namely $\mathcal{S}_i$, is, therefore, the set of all the possible non-empty subsets of T_i, where T_i is the set containing all the existing edges that connect a_i to the tasks. In ECMS, this means that an agent can report only a capacity that is lower than its real one and that it cannot report an edge that does not exist. In this case, the set of strategies of agent a_i is then $\mathcal{S}_i \times [b_i]$, where b_i is the real agent's capacity.

For the sake of simplicity, **from now on all the results and definitions we introduce are for the EMS**, unless we specify otherwise. We generalize our results to the ECMS in Section 7.

The Mechanisms. A mechanism for the MVbM problem is a function $\mathbb{M}$ that takes as input the private information of the agents and returns a $\boldsymbol{b}$-matching. We denote with $\mathcal{I}_{\mathbb{M}}$ the set of possible inputs for $\mathbb{M}$ in the EMS. In our paper, we consider three mechanisms:

(i) $\mathbb{M}_{BFS}$, which takes in input the edges of the agents and uses Algorithm 1 endowed with the BFS to select a matching[3].
(ii) $\mathbb{M}_{DFS}$, which takes in input the edges of the agents and uses Algorithm 1 endowed with the DFS to select a matching.
(iii) $\mathbb{M}_{AP}$, which takes in input the edges of the agents and uses the approximated version of Algorithm 1 to select a matching.

Notice that $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ are optimal since they are both induced by Algorithm 1.

Given a mechanism $\mathbb{M}$, every element $I \in \mathcal{I}_{\mathbb{M}}$ is composed by the reports of n self-interested agents, so that $\mathcal{I}_{\mathbb{M}} = \otimes_{i=1}^{n} \mathcal{I}_i$, where $\mathcal{I}_i$ is the set of feasible inputs for agent a_i. We say that a mechanism $\mathbb{M}$ is *strategyproof* (or, equivalently, *truthful*) for the EMS if no agent can get a higher payoff by hiding edges. More formally, if I_i is the true type of agent a_i, it holds true that

$$w_i(\mathbb{M}(I_i', I_{-i})) \leq w_i(\mathbb{M}(I_i, I_{-i}))$$

for every $I_i' \in \mathcal{S}_i$. Another important property for mechanisms is the group strategyproofness. A mechanism is *group strategyproof* for agent manipulations if no group of agents can collude to hide some of their edges in such a way that (i) the utility obtained by every agent of the group after hiding the edges is greater than or equal to the one they get by reporting truthfully, (ii) at least one agent gets a better payoff after the group hides the edges.

[3] Throughout the paper, we use $\mathbb{M}_{BFS}$ to denote both the mechanism that takes as input the edges and both the edges and the capacity of every agent. It will be clear from the context which is the input of the mechanism. The same goes for the other two mechanisms.

To evaluate the performances of the mechanisms, we use the Price of Anarchy (PoA), Price of Stability (PoS), and approximation ratio (ar), which we briefly recall in the following.

Price of Anarchy. The Price of Anarchy (PoA) of mechanism $\mathbb{M}$ is defined as the maximum ratio between the optimal social welfare and the welfare in the worst Nash Equilibrium, hence

$$PoA(\mathbb{M}) := \sup_{I \in \mathcal{I}} \frac{w(\mu(I))}{w(\mu_{wNE}(I))},$$

where $\mu_{wNE}(I)$ is the output of $\mathbb{M}$ when the agents act according to the worst Nash Equilibrium, i.e. the Nash Equilibrium that achieves the worst social welfare.

Price of Stability. The Price of Stability (PoS) of a mechanism $\mathbb{M}$ is defined as the maximum ratio between the optimal social welfare and the welfare in the best Nash Equilibrium, hence

$$PoS(\mathbb{M}) := \sup_{I \in \mathcal{I}} \frac{w(\mu(I))}{w(\mu_{bNE}(I))},$$

where $\mu_{bNE}(I)$ is the output of $\mathbb{M}$ when the agents act according to the best Nash Equilibrium, i.e. the Nash Equilibrium that achieves the maximum social welfare. Notice that, by definition, we have $PoS(\mathbb{M}) \leq PoA(\mathbb{M})$.

Approximation Ratio. The approximation ratio of a truthful mechanism $\mathbb{M}$ is defined as the maximum ratio between the optimal social welfare and the welfare returned by $\mathbb{M}$. Hence, we have

$$ar(\mathbb{M}) := \sup_{I \in \mathcal{I}} \frac{w(\mu(I))}{w(\mu_{\mathbb{M}}(I))},$$

where $\mu_{\mathbb{M}}(I)$ is the output of $\mathbb{M}$ when I is given in input.

4 The Truthfulness of the Mechanisms

In this section, we study the truthfulness of the three mechanisms induced by Algorithm 1 and its approximation version. We show that, although $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ are optimal, they are not truthful due to an impossibility result. Furthermore, we show that the manipulability of a mechanism is related to the length of the augmenting paths found during the routine of Algorithm 1 and use this characterization to prove that $\mathbb{M}_{AP}$ is truthful.

Theorem 1 *There is no deterministic truthful mechanism that always returns an MVbM.*

Proof. We show this using a counterexample. Consider two agents a_1 and a_2 and three tasks t_1, t_2, t_3. The edge set is $E = \{(a_1,t_1),(a_1,t_2),(a_2,t_1),(a_2,t_3)\}$. The values of the three tasks are $q_1 = 1, q_2 = 0.1$, and $q_3 = 0.1$, respectively, while the capacity of both agents is $b_1 = b_2 = 1$. It is easy to see that the optimal matching is not unique. In particular, both $\{(a_1,t_1), (a_2,t_3)\}$ and $\{(a_1,t_2),(a_2,t_1)\}$ are feasible solutions whose total weight is 1.1. Let us assume that the mechanism returns $\{(a_1,t_1), (a_2,t_3)\}$. In this case, if agent a_2 does not report edge (a_2,t_3), the maximum matching becomes $\{(a_1,t_2),(a_2,t_1)\}$. According to this, agent a_2's utility increases from 0.1 to 1. Similarly, if the mechanism returns $\{(a_1,t_2),(a_2,t_1)\}$, agent a_1 can manipulate the result by hiding the edge (a_1,t_2). Therefore, there is no deterministic truthful mechanism that always returns a maximum matching. □

From Theorem 1, we infer that both $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ are not truthful and, thus, not group strategyproof. In the following, we characterize a sufficient condition under which agents' best strategy is to report truthfully. This characterization will allow us to deduce that $\mathbb{M}_{AP}$ is truthful and to present sets of instances on which $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ are truthful.

Lemma 1 *Let us consider an instance in which Algorithm 1 completes its routine using only augmenting paths of length equal to* 1*. If an agent cannot misreport edges in such a way that the algorithm will find an augmenting path of length greater than* 1 *in its implementation, then the agent's best strategy is to report truthfully.*

Proof. Let us denote with μ_k the matching found at the k-th step. It is easy to see that Algorithm 1 concludes its routine using only augmenting paths of length equal to 1 if and only if the sequence of matching that it finds is monotone increasing. That is, $\mu_k \subset \mu_{k+1}$ for every k. Since the sequence $\{\mu_k\}_k$ is increasing, the matching μ_{k+1} is defined by adding (at most) an edge to μ_k. Let us now assume that an agent hides a set of edges in such a way that the matching sequence found by the algorithm is still monotone increasing. Thus, after the manipulation, the final matching is still obtained by adding (at most) an edge at every iteration. By the definition of $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$, the edge added at step j, is taken from the ones that are connected to the task t_j. Therefore, by hiding one or more edges, the manipulative agent can only reduce the total value of the tasks assigned to it, which concludes the proof. □

Since $\mathbb{M}_{AP}$ uses only augmenting paths of length at most equal to 1, regardless of the input, Lemma 1 allows us to conclude that $\mathbb{M}_{AP}$ is truthful. Moreover, due to the results proven in [19], we can characterize the approximation ration of $\mathbb{M}_{AP}$.

Theorem 2 *The mechanism $\mathbb{M}_{AP}$ is truthful with respect to agent manipulation. Moreover, its approximation ratio is* 2*.*

The approximation ratio achieved by $\mathbb{M}_{AP}$ is actually the best possible ratio achievable, as the next result shows.

Theorem 3 *There does not exist a truthful and deterministic mechanism for the MVbM problem that achieves an approximation ratio better than* 2*.*

Proof. Toward a contradiction, assume that there exists a mechanism $\mathbb{M}$ whose approximation ratio is equal to $2-\delta$, where δ is a positive value. Let us now consider the following instance. We have two agents, namely a_1 and a_2, and two tasks, namely t_1 and t_2. The capacity of both agents is set to be equal to 1. The value q_1 of t_1 is $1+\epsilon$ and the value q_2 of t_2 is equal to 1. Finally, we assume that, according to their truthful inputs, both agents are connected to both tasks, so that $E=\{(a_1,t_1),(a_1,t_2),(a_2,t_1),(a_2,t_2)\}$. It is easy to see that the maximum value that a matching can achieve is $2+\epsilon$. If the mechanism $\mathbb{M}$ does not allocate both the tasks, we have that the value achieved by the mechanism is, at most $1+\epsilon$. Therefore, we have that the approximated ratio of $\mathbb{M}$ is at least equal to $\frac{2+\epsilon}{1+\epsilon}=1+\frac{1}{1+\epsilon}$. If we take $\epsilon<\frac{\delta}{1-\delta}$, we get that the approximated ratio of $\mathbb{M}$ should be greater than $2-\delta$, which is a contradiction. Hence $\mathbb{M}$ allocates both the tasks in the previously described instance. Without loss of generality, let us assume that $\mathbb{M}$ allocates t_1 to a_2 and t_2 to a_1. Let us now consider the instance in which a_1 is not connected with the task t_2, so that $E'=\{(a_1,t_1),(a_2,t_1),(a_2,t_2)\}$. The maximal value that a matching can achieve is still $2+\epsilon$. However, since the mechanism is truthful, the agent a_1 cannot receive any task. Indeed, the only task that $\mathbb{M}$ can assign to a_1 is t_1, however, if $\mathbb{M}$ assigns t_1 to a_1, it means that agent a_1 can manipulate $\mathbb{M}$ by reporting the set of edges $\{(a_1,t_1)\}$ over $\{(a_1,t_1),(a_1,t_2)\}$ in the instance when the truthful input is E. Since the first agent cannot receive any tasks from $\mathbb{M}$, we have that the maximum matching value achieved by $\mathbb{M}$ when the input is E' is, at most, $1+\epsilon$, so that the approximation ratio of $\mathbb{M}$ is, at least $\frac{2+\epsilon}{1+\epsilon}$. Again, by taking $\epsilon<\frac{\delta}{1-\delta}$, we conclude a contradiction. □

From Theorem 3, we then infer that $\mathbb{M}_{AP}$ is the best possible deterministic and truthful mechanism for our game theoretical setting. To conclude, we show that, $\mathbb{M}_{AP}$ is also group strategyproof if all the tasks have different values.

Theorem 4 *If all the tasks in T have different values, then $\mathbb{M}_{AP}$ is group strategyproof.*

Proof. Toward a contradiction, let us assume that there exists a coalition of agents $C=\{a_{i_1},\dots,a_{i_\ell}\}$ that is able to collude. Since hiding edges that are not returned by $\mathbb{M}_{AP}$ does not alter the outcome of the mechanism, we assume, without loss of generality, that at least agent a_{i_1}, hides one of the edges that are in the matching found by $\mathbb{M}_{AP}$ when all the agents report truthfully. Let us denote with t_l the task connected to a_{i_1} through the hidden edge. After misreporting agent a_{i_1} cannot be allocated with t_l. Furthermore, due to the routine of $\mathbb{M}_{AP}$, each task is allocated independently from the others, hence a_{i_1} cannot be allocated with a better task. Finally, since there are no tasks with the same value, a_{i_1}'s payoff is necessarily lowered by misreporting, even if in a coalition, which is a contradiction. □

The condition of Theorem 4 are tight. Indeed, as we show in the next example, even if just two tasks have the same value, the mechanism is no longer group strategyproof.

Example 1 *Let us consider the following instance. The set of agents is composed of three elements, namely a_1, a_2, and a_3. The capacity of each agent is set to* 1*, so that $b_1=b_2=b_3=1$. The set of tasks is composed of two elements, namely t_1 and t_2. The value of both tasks is equal to* 1*. Finally, let $E=\{(a_1,t_1),(a_1,t_2),(a_2,t_2),(a_3,t_1)\}$ be the truthful input. Then, $\mathbb{M}_{AP}(E)=\{(a_1,t_1),(a_2,t_2)\}$. However a_1 and a_3 can collude: if agent a_1 hides the edge (a_1,t_1), the $\mathbb{M}_{AP}$ returns $\mu'=\{(a_1,t_2),(a_3,t_1)\}$.*

5 The Price of Anarchy and the Price of Stability

In this section, we study to what extent an agent can manipulate $\mathbb{M}_{BFS}$ or $\mathbb{M}_{DFS}$ in its favor. First, we show that an agent who is not matched to any task when reporting truthfully cannot improve its utility by manipulation.

Lemma 2 *Given a truthful input, if an agent receives a null utility from $\mathbb{M}_{BFS}$, then its utility cannot be improved by hiding edges. The same holds for the mechanism $\mathbb{M}_{DFS}$.*

The latter Lemma ensures us that, without loss of generality, if an agent is able to manipulate $\mathbb{M}_{BFS}$ or $\mathbb{M}_{DFS}$, then the same agent is assigned at least one task when it reports truthfully. We now show that the first agent processed by either $\mathbb{M}_{BFS}$ or $\mathbb{M}_{DFS}$, hence a_1, is always able to get its highest possible payoff by misreporting.

Theorem 5 *For both $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$, agent a_1's best strategy is to report only the top b_1-valued tasks to which it is connected.*

Proof. Let us denote with $t_{j_1}, \ldots, t_{j_{b_1}}$ the top b_1 tasks agent a_1 is connected to. For every t_{j_r}, at the j_r-th step of Algorithm 1, agent a_1 will not be saturated; therefore the path $P_{j_r} = \{(t_{j_r}, a_1)\}$ is augmenting with respect to matching μ_{j_r-1}. Moreover, since the BFS searches among the vertices in lexicographical order, the path P_{j_r} is always the first one being explored and, since it is augmenting, it will be the one returned. To conclude, we notice that, after the j_r-th iteration, there are no augmenting paths that pass by a_1 as all the edges connected to a_1 are already in the matching, thus the set of tasks allocated to a_1 will not change in later iterations of the algorithm. By a similar argument, we infer the same conclusion for $\mathbb{M}_{DFS}$. □

In the next example, we show that the advantage described in Theorem 5 is only due to the fact that agent a_1 is self-aware of its position in the processing process.

Example 2 *Let us consider the following instance. There are* 3 *agents, namely* α*,* β*, and* γ *whose capacities are* $b_\alpha = 2$ *and* $b_\beta = b_\gamma = 1$*. Let us consider a set of* 4 *tasks, namely* t_j *with* $j \in [4]$ *whose respective values are* $q_j = 2^{-j}$*. In the truthful input, agent* α *is connected to all* 4 *tasks, while agent* β *is connected only to task* t_1 *and agent* γ *is connected only to task* t_2*. Let us consider the mechanism* $\mathbb{M}_{BFS}$*: the maximum matching for the truthful input allocates* t_1 *to agent* β*,* t_2 *to agent* γ*, and the other two tasks to agent* α*. Let us now assume that the processing order of the agents is* α*,* β*, and* γ*. That is,* $a_1 = \alpha$*,* $a_2 = \beta$*, and* $a_3 = \gamma$*. Then, if agent* α *applies the strategy highlighted in Theorem 5 and reports only the first two edges, it improves its utility from* $q_3 + q_4$ *to* $q_1 + q_2$*. However, in a different order of agents in which agent* α *is the second, i.e.,* $a_2 = \alpha$*, if it applies the same strategy, then agent* α *gets only one of the two tasks (depending on who is the agent going first) while if it goes as the third, it receives no tasks. We also note that, if agent* α *is the second, its best strategy is to report the edges connecting it to tasks* t_1 *and* t_3 *if agent* γ *goes first or tasks* t_2 *and* t_3 *if* β *goes first. Hence, the priority of agent* α *and the priority of the other two agents determines what the best strategy for agent* α *is. Similarly, the same conclusion can be drawn for the mechanism* $\mathbb{M}_{DFS}$*.*

In general, the best strategy of an agent depends on its priority and the reports of the agents whose priority is higher than its. Indeed, once the agents' order is fixed, it is possible to describe the Nash Equilibria of both $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$. For every $i = 0, 1, \ldots, n$, let us define the sets $T^{(i)}$ and $B^{(i)}$ in the following iterative way: (i) $T^{(0)} = T$ where T is the set of all the tasks and $B^{(0)} = \emptyset$; (ii) $T^{(i)} = T^{(i-1)} \backslash B_i^{(i-1)}$, where $B_i^{(i-1)}$ is the set containing the top b_i-valued tasks among the ones in $T^{(i-1)}$ that agent i is connected to. If agent i is connected to less than b_i tasks, then $B_i^{(i-1)}$ contains all the tasks in $T^{(i-1)}$ to which agent i is connected to. We then define agent a_i's *First-Come-First-Served* (FCFS) policy as $FCFSP_i = \{(a_i, t_j)\}_{t_j \in B_i^{(i-1)}}$. Notice that $FCFSP_i \in \mathcal{S}_i$ for every $i \in [n]$, so that it is a feasible strategy for every agent.

Theorem 6 *Given an MVbM problem, the FCFS policies constitute a Nash Equilibrium that achieves the lowest social welfare for both mechanisms* $\mathbb{M}_{BFS}$ *and* $\mathbb{M}_{DFS}$*. Moreover, we have*

$$\mathbb{M}_{BFS}(\cup_{a_i \in A} FCFSP_i) = \mathbb{M}_{DFS}(\cup_{a_i \in A} FCFSP_i) = \cup_{a_i \in A} FCFSP_i. \quad (1)$$

Proof. We prove the first part of the theorem in two steps. First, we prove that the FCFS policies constitute a Nash Equilibrium. Second, we show that the Nash Equilibrium they form is the one with the lowest possible social welfare.

Let us then consider agent a_i and, toward a contradiction, let us assume that, when all the other agents report their FCFS policy, reporting the set of edges $S_{a_i} \neq FCFSP_i$ gives a_i a bigger payoff than what it would get from reporting $FCFSP_i$. Let us set $s_i = \min_{(a_i, t_j) \in FCFSP_i} q_j$. Let us assume that S_{a_i} contains an edge that connects a_i to a task that has a higher value than s_i, namely t_l. We now show that, since the other agents are applying their FCFS policies, the task t_l is allocated to an agent with higher priority unless the edge already belongs to $FCFSP_i$. Indeed, since $|FCFSP_j| \leq b_j$ for every $a_j \in A$, there cannot be augmenting paths that pass by any of the agents playing their FCFS policy. Indeed, since the union of the FCFS policies is a b-matching, either $(a_i, t_l) \in FCFSP_i$ or there exists another agent whose FCFS policy connects it to t_l. If t_l is connected to an agent a_k, $(a_k, t_l) \in FCFSP_k$, and agent a_k's priority is higher than agent a_i's priority, the final output of the mechanism assigns t_l to a_k. We can then assume that S_{a_i} does not contain edges that connect agent a_i to tasks with values higher than s_i and that do not belong to $FCFSP_i$. To conclude, we notice that if S_{a_i} contains an edge connecting agent a_i to a task that has a value lower than s_i, then the payoff of agent a_i can only be lowered. Indeed, if $|FCFSP_i| = b_i$, then a_i cannot improve its payoff by reporting edges that connect a_i to tasks that have a value lower than s_i. This follows from the fact that all the other players are using their FCFS policies and therefore agent a_i is allocated the set $B_i^{(i-1)} = \{t_j \in T \text{ s.t. } (a_i, t_j) \in FCFSP_i\}$ if it uses its FCFS policy. If $|FCFSP_i| < b_i$, by definition, it means that there are no tasks that agent a_i can be connected to and that have a value lower than s_i. Therefore S_{a_i} does not contain edges connecting a_i to a task with a value lower than s_i nor edges connecting it with tasks that have a value greater than s_i and that are not included in $FCFSP_i$. Since reporting a subset of $FCFSP_i$ would result in a lower payoff, we deduce that $S_{a_i} = FCFSP_i$, which is a contradiction since we assumed $FCFSP_i \neq S_{a_i}$.

We now prove that the Nash Equilibrium induced by the FCFS policies is one of the worst equilibria. Toward a contradiction, let us consider another set of strategies, namely $\{S_{a_i}\}_{a_i \in A}$, such that the social welfare achieved by this equilibrium is strictly lower than the one obtained if every agent uses its FCFS policy. By the definition of social welfare, there must exist at least one agent that, according to the equilibrium defined by $\{S_{a_i}\}_{a_i \in A}$, receives a payoff that is strictly lower than the one it would obtain by using its FCFS policy. Let us denote with a_k the first agent that, according to the priority order of the mechanism, receives a lower value. Agent a_k cannot be the first agent, as it otherwise could apply its FCFS policy and get a better payoff. Then, agent a_k is among the remaining agents and it is not getting any of the tasks that are given to the first agent according to its FCFS policy, as otherwise, agent a_1 could increase its payoff by manipulating and the set of strategies $\{S_i\}_{a_i \in A}$ would not be a Nash Equilibrium. From a similar argument, we infer that agent a_k cannot be the second agent, as otherwise, it could get a better payoff by using its FCFS policy. Moreover, the second agent is allocated the tasks that are granted to it from its FCFS policy. Both of which we have already proved cannot be. By applying the same argument to the other agents, we get a contradiction, since no agent can be agent a_k. We, therefore, conclude that the set of strategies given by the FCFS policies is one of the worst Nash Equilibrium.

The last part of the Theorem follows from the fact that $\cup_{a_i \in A} FCFSP_i$ is itself a b-matching, hence it is also an $MVbM$ with respect to the edge set $E = \cup_{a_i \in A} FCFSP_i$. □

Following the same argument presented in the proof of Theorem 6, we infer that any set of strategies $\{S_i\}_{a_i \in A}$ for which it holds $FCFSP_i \subset S_i$ and $\min_{t_j,\,(a_i,t_j)\in FCFSP_i} q_j = \min_{t_j,\,(a_i,t_j)\in S_i} q_j$ for every i, defines a Nash Equilibrium. Among this class of Nash Equilibria, $\{FCFSP_i\}_{a_i\in A}$ is the only one for which the identities in (1) hold. Moreover, this equilibrium achieves the same social welfare of the matching returned by $\mathbb{M}_{AP}$.

Theorem 7 *Given an MVbM problem, let μ_E be the matching returned by $\mathbb{M}_{AP}$. Then, it holds that $\mu_E = \cup_{a_i\in A} FCFSP_i$. That is, the output of $\mathbb{M}_{AP}$ on any given instance is equal to the union of the agents' FCFS policies. Hence, the social welfare achieved by $\mathbb{M}_{AP}$ is equal to the social welfare achieved by $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ in one of their worst Nash Equilibria.*

Proof. We prove this Theorem by induction. First, we prove that $\mathbb{M}_{AP}$ allocates a_1 with the set of tasks $FCFSP_1$. Second, we show that if all the agents $a_1, \dots, a_k$ receive $FCFSP_1, \dots, FCFSP_k$ from $\mathbb{M}_{AP}$, then also agent a_{k+1} receives $FCFSP_{k+1}$.

Let us consider a_1. Since $\mathbb{M}_{AP}$ checks the agents following their orders, a_1 is always the first agent checked. Hence, a task t_j is not allocated to a_1 if and only if there are b_1 other tasks that have a higher value than t_j and that are all connected to a_1. This means that the set of tasks allocated to a_1 is $FCFSP_1$.

Let us now assume that agents $a_1, a_2, \dots, a_k$ are given the tasks contained in $FCFSP_1, FCFSP_2, \dots, FCFSP_k$ respectively, and let us consider a_{k+1}. We have that $\mathbb{M}_{AP}$ allocates a task t_j to a_{k+1} if, at step j of Algorithm 1, all the agents with priorities higher than a_{k+1} that are connected to t_j are already saturated and a_{k+1} is not saturated. However, by assumption, agents with a higher priority than a_{k+1} are getting the tasks that they would get from their FCFS policies. We, therefore, conclude that the tasks allocated to agent a_{k+1} from $\mathbb{M}_{AP}$ consist of a subset of $T^{(k)}$. From an argument similar to the one used for agent a_1, we infer that a_{k+1} receives the top b_{k+1} higher valued tasks among the ones in $T^{(k)}$, which coincides with the set $FCFSP_{k+1}$. We conclude the proof by induction. □

Finally, we show that Theorem 7 along with Theorem 2 allows us to compute the PoA of both $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$.

Theorem 8 *The PoA of $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ is equal to 2.*

Proof. From Theorem 7, the matching returned by $\mathbb{M}_{AP}$ achieves a social welfare equal to the social welfare of one of the worst Nash Equilibrium of $\mathbb{M}_{BFS}$. Since $\mathbb{M}_{BFS}$ returns an MVbM, we have

$$PoA(\mathbb{M}_{BFS}) = \sup_{I\in\mathcal{I}} \frac{w(\mu(I))}{w(\mu_{wNE}(I))} = \sup_{I\in\mathcal{I}} \frac{w(\mathbb{M}_{BFS}(I))}{w(\mathbb{M}_{AP}(I))}. \quad (2)$$

Since the matching found by $\mathbb{M}_{BFS}$ achieves the maximum social welfare, the last term in equation (2) is bounded from above by the $a.r.(\mathbb{M}_{AP})$, so that $PoA(\mathbb{M}_{BFS}) \leq ar(\mathbb{M}_{AP}) = 2$. To conclude $PoA(\mathbb{M}_{BFS}) = 2$ we show a lower bound of 2. The set of agents is composed of two agents, namely a_1 and a_2, we assume the agents to be ordered according to the algorithm priority. The capacity of both agents is equal to 1. The set of tasks contains two tasks, namely t_1 and t_2, whose values are $1+\epsilon$ and 1, respectively. Finally, let us assume that the truthful input is given by $E = \{(a_1,t_1),(a_1,t_2),(a_2,t_1)\}$. The social welfare is then equal to $2+\epsilon$. However, in the worst Nash Equilibrium, the welfare is $1+\epsilon$. By taking the limit for $\epsilon \to 0$, we conclude that the PoA is equal to 2. By a similar argument, we infer $PoA(\mathbb{M}_{DFS}) = 2$. □

The previous bound is tight: there does not exist a deterministic mechanism for the $MVbM$ problem that has a PoA lower than 2.

Theorem 9 *For every deterministic mechanism $\mathbb{M}$, we have $PoA(\mathbb{M}) \geq 2$ with respect to agent manipulations.*

Proof. Toward a contradiction, let $\mathbb{M}$ be a deterministic mechanism whose PoA is lower than 2. Let us consider the following instance. We have two tasks, namely t_1 and t_2, whose values are $1+\epsilon$ and 1, respectively. We then have two agents, namely a_1 and a_2 and both have a capacity equal to 1. Let us now consider the instance in which both the agents are only connected to task t_1, hence the truthful input is $E = \{(a_1,t_1),(a_2,t_1)\}$.

Since $PoA(\mathbb{M})$ is finite, we have that $\mathbb{M}$ allocates t_1 to one agent. Indeed, if no agent receives a task, no one can improve its own payoff by hiding its only edge (we recall that each agent is bounded by its statements). Hence, the truthful instance is already a Nash Equilibrium. Furthermore, since the social welfare of this Nash Equilibrium is 0, this is also one of the worst Nash Equilibria, thus we find a contradiction since we assumed that $\mathbb{M}$ has a finite PoA.

Let us then assume that one agent gets t_1. Without loss of generality, let us assume that $\mathbb{M}$ allocates t_1 to a_1, the other case is completely symmetric with respect to the one we are about to present.

Let us now consider the instance whose truthful input is $E = \{(a_1,t_1),(a_1,t_2),(a_2,t_1)\}$. If $\mathbb{M}$ allocates t_1 to a_2, we have that a Nash Equilibrium is obtained when agent a_1 hides arc (a_1,t_2). Indeed, if agent a_1 hides (a_1,t_2), the input of the mechanism is $E = \{(a_1,t_1),(a_2,t_1)\}$ which gives the first task to a_1. Since a_2 is bounded by its statements and its only alternative is to report no edges, it has no better strategy to play. Similarly, since a_1 is getting its best possible payoff, it has no better strategy to play. We then conclude that when a_1 hides the edge (a_1,t_2), we have a Nash Equilibrium. Finally, we observe that, by taking ϵ small enough, we get a contradiction with the assumption $PoA(\mathbb{M}) < 2$, since the maximum social welfare is $2+\epsilon$, while the social welfare returned by the mechanism in the worst Nash Equilibrium is at most $1+\epsilon$. Notice that there might be another Nash Equilibrium in which the social welfare is even lower, however, it suffice to notice that the social welfare achieved in the worst Nash Equilibrium is lower than $1+\epsilon$ to conclude the proof. Similarly, if the mechanism does not allocate t_1 to a_2, the instance is already a Nash Equilibrium. Indeed, by the same argument used before, a_2 cannot improve its own payoff, since it is getting no tasks. If a_1 is allocated with the task, it cannot improve its payoff either, since it is getting the maximum payoff it can get. Finally, if a_1 is not getting t_1, it can hide (a_1,t_2) and return to the instance we considered before. Again, by taking ϵ small enough, we retrieve that the $PoA(\mathbb{M})$ cannot be less than 2. □

We close the section by studying the PoS of $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$. We recall that the PoA of every mechanism is greater than its PoS, thus we infer that both $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ have a PoS at most equal to 2. Moreover, since the Nash Equilibrium in the example we used in the proof of Theorem 9 is unique, the best and worst Nash Equilibria achieve the same social welfare. In particular, this allows us to prove that the PoS of both $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ is equal to 2. Furthermore, this value is tight for the class of deterministic mechanisms.

Theorem 10 *The PoS of $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ is equal to 2. Moreover, no deterministic mechanism can achieve a lower PoS.*

6 The Truthful Inputs for $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$

In this section, we describe three sets of inputs in which $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ are truthful. In particular, we consider the three following settings. In Theorem 11, we consider the case in which there is a shortage of tasks, so that the capacity of each agent exceeds the number of tasks to which it is connected. In particular, no agent can be saturated. In Theorem 12, instead, we describe what happens when every task can be contended by another agent. In Theorem 13, we study the case in which the private information of the agents can be clustered, that is different agents are connected to the same set of tasks and the same capacity. Going back to the worker-project example at the beginning of the paper, this means that the connection between the worker and the project depends, for example, on the field of expertise of the worker or their background formation.

Theorem 11 *Let us consider the set of inputs such that, according to E, all the agents' degrees are less than or equal to their capacity, that is $\sum_{t_j \in T_i} e_{i,j} \leq b_i$ for every $i \in [n]$, then $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ are truthful on this set of inputs.*

Proof. Since no agents can be saturated, any augmenting path found by the mechanisms has length equal to 1. Moreover, since hiding edges cannot lead an agent to be saturated, Lemma 1 allows us to conclude the proof. □

This is the only class of inputs we are considering on which $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ behave in the same way. In the other two frameworks, $\mathbb{M}_{DFS}$ is not truthful, while its counterpart $\mathbb{M}_{BFS}$ is. This is due to the fact that BFS searches for the shortest possible augmenting path in its execution.

Theorem 12 *Let μ be the matching returned by $\mathbb{M}_{BFS}$. If for every task t_j there exists an edge $e \notin \mu$ that connects t_j to an unsaturated agent, then no agents can increase their utility by hiding only one edge. In particular, if the truthful input is a complete bipartite graph and the vector of the capacities $\boldsymbol{b}$ is such that $m \leq \sum_{i=1}^{n} b_i - \max_{i \in [n]} b_i$, then the best strategy for every agent is to report truthfully.*

Proof. Assume, toward a contradiction, that an agent, namely a_i gets a benefit by hiding an edge, namely $e = (a_i, t_j)$. By hypothesis, we have that every task is connected to an unsaturated agent. Let a_k be one of the unsaturated agents to which t_j is connected to. Then BFS will always find an augmenting path whose length is 1 when it is asked to allocate t_j, since there exists the augmenting path (a_k, t_j). Therefore, by Lemma 1 we infer a contradiction. □

Let $\mathcal{A} = \{A^{(1)}, \dots, A^{(r)}\}$ be a partition of A. We say that $A^{(\ell)}$ is the ℓ-th class of the agents. Since $\mathcal{A}$ is a partition, every agent $a_i \in A$ belongs to only one class. Let us assume that the capacity b_i and the set of edges T_i of every agent $a_i \in A$ depends only on the class $A^{(\ell)}$ to which a_i belongs, so that $b_i = b^{(\ell)}$ and $T_i = T^{(\ell)}$. Using an argument that is similar (but more delicate) to the one used in Theorem 12, we are able to prove that $\mathbb{M}_{BFS}$ is truthful if every class contains enough agents.

Theorem 13 *In the framework described above, if $|A^{(\ell)}| > \lceil \frac{|T^{(\ell)}|}{b^{(\ell)}} \rceil + 1$, then no agents belonging to the ℓ-th class can benefit by misreporting to $\mathbb{M}_{BFS}$.*

In the appendix, we report two examples that show that both Theorem 12 and Theorem 13 do not hold for $\mathbb{M}_{DFS}$.

7 Agents Manipulating their Edges and Capacity

Finally, we extend our study on the truthfulness of $\mathbb{M}_{BFS}$, $\mathbb{M}_{DFS}$, and $\mathbb{M}_{AP}$ to the ECMS, i.e. we allow the agents self-report their capacity along with their edges. As for the EMS, we assume the agents to be bounded by their statements, thus they can manipulate only by hiding edges or by reporting a lower capacity than their real one. In this setting, a mechanism $\mathbb{M}$ is truthful if, for every $i \in [n]$, it holds $w_i((I'_i, b'_i), J_{-i}) \leq w_i((I_i, b_i), J_{-i})$, for every $(I'_i, b'_i) \in \mathcal{S}_i \times [b_i]$, where J_{-i} are the reports of the other agents. Once we fix the set of strategies of each agent, we can define the PoA, PoS, and approximation ratio of a mechanism $\mathbb{M}$ as for the EMS by carefully changing the set of strategies to fit the ECMS case.

With a slight abuse of notation, we still use $\mathbb{M}_{BFS}$, $\mathbb{M}_{DFS}$, and $\mathbb{M}_{AP}$ to denote the mechanisms obtained from Algorithm 1 and its approximation version. As we show, neither the truthfulness nor the efficiency guarantees of $\mathbb{M}_{BFS}$, $\mathbb{M}_{DFS}$, and $\mathbb{M}_{AP}$ change from EMS to ECMS. Furthermore, all the bounds are still tight.

Theorem 14 *In the ECMS, $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ are both not truthful. The PoA and the PoS of both $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$ are equal to 2. Moreover, these bounds are tight, hence there is no other deterministic mechanism whose PoA or PoS is lower.*

Similarly, $\mathbb{M}_{AP}$ is still truthful in the ECMS, and its approximation ratio is unchanged.

Theorem 15 *In the ECMS, $\mathbb{M}_{AP}$ is truthful, its approximation ratio is equal to 2. Moreover, there is no deterministic truthful mechanism with a better approximation ratio. Finally, if the tasks have different values, $\mathbb{M}_{AP}$ is group strategyproof.*

8 Conclusion and Future Work

In this paper, we propose a new game-theoretical framework for MVbM problems, where one side of the bipartite graph consists of agents and the other side consists of tasks with objective values. We consider scenarios where agents can behave strategically by hiding connections with tasks or lowering their capacity. We analyze three mechanisms in this framework: $\mathbb{M}_{BFS}$, $\mathbb{M}_{DFS}$, and $\mathbb{M}_{AP}$. We first show that these mechanisms are either optimal ($\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$) or truthful ($\mathbb{M}_{AP}$). Then, we demonstrate that these mechanisms are also the best in terms of PoA (Price of Anarchy), PoS (Price of Stability), and approximation ratio. In other words, no other mechanisms can outperform these ones with respect to these performance measures in our setting.

A future direction of our work is to study the effect of shuffling the agents' order on the manipulability of $\mathbb{M}_{BFS}$. Specifically, we are interested in investigating whether randomizing the priority of agents further highlights the differences between $\mathbb{M}_{BFS}$ and $\mathbb{M}_{DFS}$. Moreover, it would be interesting to explore how placing bounds on the number of edges that each agent can report affects the performance and manipulability of the mechanisms we studied. Finally, we plan to investigate more in details other classic game-theoretical aspects of our model, such as fairness and envy-freeness.

Acknowledgements

We would like to thank the referees for their comments, which helped improve this paper considerably. This project is partially supported by a Leverhulme Trust Research Project Grant (2021 – 2024). Jie Zhang is also supported by the EPSRC grant (EP/W014912/1).

References

[1] Atila Abdulkadiroğlu, 'College admissions with affirmative action', *International Journal of Game Theory*, **33**(4), 535–549, (2005).

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

[3] Atila Abdulkadiroğlu, Yeon-Koo Che, Parag A Pathak, Alvin E Roth, and Olivier Tercieux, 'Efficiency, Justified Envy, and Incentives in Priority-Based Matching', *American Economic Review: Insights*, **2**(4), 425–442, (2020).

[4] Ahmed Al-Herz and Alex Pothen, 'A 2/3-approximation algorithm for vertex-weighted matching', *Discrete Applied Mathematics*, **308**, 46–67, (2022).

[5] Gennaro Auricchio, Federico Bassetti, Stefano Gualandi, and Marco Veneroni, 'Computing Wasserstein Barycenters via Linear Programming', in *Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pp. 355–363, Cham, (2019). Springer International Publishing.

[6] Yossi Azar, Martin Hoefer, Idan Maor, Rebecca Reiffenhäuser, and Berthold Vöcking, 'Truthful Mechanism Design via Correlated Tree Rounding', in *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pp. 415–432, (2015).

[7] Haris Aziz, Péter Biró, and Makoto Yokoo, 'Matching Market Design with Constraints', in *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 12308–12316, (2022).

[8] Haris Aziz, Serge Gaspers, Zhaohong Sun, and Toby Walsh, 'From Matching with Diversity Constraints to Matching with Regional Quotas', in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '19, p. 377–385, Richland, SC, (2019). International Foundation for Autonomous Agents and Multiagent Systems.

[9] Michel Balinski and Tayfun Sönmez, 'A Tale of Two Mechanisms: Student Placement', *Journal of Economic theory*, **84**(1), 73–94, (1999).

[10] M. Bayati, D. Shah, and M. Sharma, 'Maximum weight matching via max-product belief propagation', in *Proceedings. International Symposium on Information Theory, 2005. ISIT 2005.*, pp. 1763–1767, (2005).

[11] Benjamin E. Birnbaum and Claire Mathieu, 'On-line Bipartite Matching Made Simple', *SIGACT News*, **39**(1), 80–87, (2008).

[12] Péter Biró, Tamás Fleiner, Robert W. Irving, and David F. Manlove, 'The College Admissions problem with lower and common quotas', *Theor. Comput. Sci.*, **411**(34-36), 3136–3153, (2010).

[13] Niclas Boehmer and Klaus Heeger, 'A Fine-Grained View on Stable Many-to-one Matching Problems with Lower and Upper Quotas', *ACM Transactions on Economics and Computation*, **10**(2), 1–53, (2022).

[14] Katarína Cechlárová, Laurent Gourvès, and Julien Lesca, 'On the Problem of Assigning PhD Grants', in *28th International Joint Conference on Artificial Intelligence (IJCAI 2019)*, pp. 130–136, (2019).

[15] Deeparnab Chakrabarty and Chaitanya Swamy, 'Welfare Maximization and Truthfulness in Mechanism Design with Ordinal Preferences', in *Proceedings of the 5th conference on Innovations in Theoretical Computer Science*, pp. 105–120, (2014).

[16] Jiehua Chen, Robert Ganian, and Thekla Hamm, 'Stable Matchings with Diversity Constraints: Affirmative Action is beyond NP', in *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, IJCAI'20, (2021).

[17] Yan Chen and Tayfun Sönmez, 'Improving Efficiency of On-Campus Housing: An Experimental Study', *American Economic Review*, **92**(5), 1669–1686, (2002).

[18] Lizhong Ding, Zheyun Feng, and Yongsheng Bai, 'Clustering analysis of microrna and mRNA expression data from TCGA using maximum edge-weighted matching algorithms', *BMC medical genomics*, **12**(1), 1–27, (2019).

[19] Florin Dobrian, Mahantesh Halappanavar, Alex Pothen, and Ahmed Al-Herz, 'A 2/3-approximation algorithm for vertex weighted matching in bipartite graphs', *SIAM Journal on Scientific Computing*, **41**(1), A566–A591, (2019).

[20] Thomas. E. Easterfield, 'A combinatorial algorithm', *Journal of The London Mathematical Society*, **21**, 219–226, (1946).

[21] Jack Edmonds and Delbert Ray Fulkerson, 'Bottleneck Extrema', *Journal of Combinatorial Theory*, **8**(3), 299–306, (1970).

[22] Lars Ehlers, Isa E Hafalir, M Bumin Yenmez, and Muhammed A Yildirim, 'School choice with controlled choice constraints: Hard bounds versus soft bounds', *Journal of Economic theory*, **153**, 648–683, (2014).

[23] Salman Fadaei and Martin Bichler, 'A Truthful Mechanism for the Generalized Assignment Problem', *ACM Transactions on Economics and Computation (TEAC)*, **5**(3), 1–18, (2017).

[24] Daquan Feng, Lu Lu, Yi Yuan-Wu, Geoffrey Ye Li, Gang Feng, and Shaoqian Li, 'Device-to-Device Communications Underlaying Cellular Networks', *IEEE Transactions on communications*, **61**(8), 3541–3551, (2013).

[25] Daniel Fragiadakis, Atsushi Iwasaki, Peter Troyan, Suguru Ueda, and Makoto Yokoo, 'Strategyproof Matching with Minimum Quotas', *ACM Transactions on Economics and Computation (TEAC)*, **4**(1), 1–40, (2016).

[26] Komei Fukuda and Tomomi Matsui, 'Finding All the Perfect Matchings in Bipartite Graphs', *Applied Mathematics Letters*, **7**(1), 15–18, (1994).

[27] David Gale and Lloyd S Shapley, 'College admissions and the stability of marriage', *The American Mathematical Monthly*, **69**(1), 9–15, (1962).

[28] Masahiro Goto, Atsushi Iwasaki, Yujiro Kawasaki, Ryoji Kurata, Yosuke Yasuda, and Makoto Yokoo, 'Strategyproof matching with regional minimum and maximum quotas', *Artificial intelligence*, **235**, 40–57, (2016).

[29] Anupam Gupta, Amit Kumar, and Cliff Stein, 'Maintaining Assignments Online: Matching, Scheduling, and Flows', in *SODA*, pp. 468–479, (2014).

[30] Frank L. Hitchcock, 'The distribution of a product from several sources to numerous localities', *Journal of Mathematics and Physics*, **20**(1-4), 224–230, (1941).

[31] Aanund Hylland and Richard Zeckhauser, 'The Efficient Allocation of Individuals to Positions', *Journal of Political Economy*, **87**(2), 293–314, (1979).

[32] Yuichiro Kamada and Fuhito Kojima, 'Recent Developments in Matching with Constraints', *American Economic Review*, **107**(5), 200–204, (2017).

[33] Yuichiro Kamada and Fuhito Kojima, 'Fair matching under constraints: Theory and applications', *Review of Economic Studies*, rdad046, (2023).

[34] L. Kantorovitch, 'On the translocation of masses', *Management Science*, **5**(1), 1–4, (1958).

[35] Fuhito Kojima, 'New Directions of Study in Matching with Constraints', in *The Future of Economic Design*, 479–482, Springer, (2019).

[36] Piotr Krysta, David Manlove, Baharak Rastegari, and Jinshan Zhang, 'Size versus Truthfulness in the House Allocation Problem', in *Proceedings of the fifteenth ACM conference on Economics and computation*, pp. 453–470, (2014).

[37] Nathaniel Lahn, Sharath Raghvendra, and Jiacheng Ye, 'A Faster Maximum Cardinality Matching Algorithm with Applications in Machine Learning', *Advances in Neural Information Processing Systems*, **34**, 16885–16898, (2021).

[38] László Lovász and Michael D Plummer, *Matching theory*, volume 367, American Mathematical Society, Michigan, USA, 2009.

[39] David Manlove, *Algorithmics of matching under preferences*, volume 2, World Scientific, Singapore, 2013.

[40] Leandro Soriano Marcolino, Albert Xin Jiang, and Milind Tambe, 'Multi-agent team formation: Diversity beats strength?', in *Twenty-Third International Joint Conference on Artificial Intelligence*. Citeseer, (2013).

[41] Aranyak Mehta, 'Online Matching and Ad Allocation', *Found. Trends Theor. Comput. Sci.*, **8**(4), 265–368, (2013).

[42] Thomas H. Spencer and Ernst W. Mayr, 'Node Weighted Matching', in *Automata, Languages and Programming*, ed., Jan Paredaens, pp. 454–464, Berlin, Heidelberg, (1984). Springer Berlin Heidelberg.

[43] Robert L. Thorndike, 'The problem of classification of personnel', *Psychometrika*, **15**, 215–235, (1950).

[44] Li Wang, Huaqing Wu, Wei Wang, and Kwang-Cheng Chen, 'Socially Enabled Wireless Networks: Resource Allocation via Bipartite Graph Matching', *IEEE Communications Magazine*, **53**(10), 128–135, (2015).

[45] Yu Yokoi, 'A Generalized Polymatroid Approach to Stable Matchings with Lower Quotas', *Mathematics of Operations Research*, **42**(1), 238–255, (2017).

[46] Chuanli Zhao and Hengyong Tang, 'Single machine scheduling with general job-dependent aging effect and maintenance activities to minimize makespan', *Applied Mathematical Modelling*, **34**(3), 837–841, (2010).