

Truthful Many-to-many Assignment with Private Weights^{*}

Bruno Escoffier¹, Jérôme Monnot¹, Fanny Pascual², and Olivier Spanjaard²

¹ LAMSADE-CNRS, UMR 7243, Université Paris Dauphine
Place du M^{aitre} de Lattre de Tassigny, F-75775 Paris Cedex 16, France
{escoffier,monnot}@lamsade.dauphine.fr

² LIP6-CNRS, UMR 7606, Université Pierre et Marie Curie (UPMC)
4 Place Jussieu, F-75252 Paris Cedex 05, France
{fanny.pascual,olivier.spanjaard}@lip6.fr

Abstract. This paper is devoted to the study of truthful mechanisms without payment for the many-to-many assignment problem. Given n agents and m tasks, a mechanism is truthful if no agent has an incentive to misreport her values on the tasks (agent a_i reports a score w_{ij} for each task t_j). The one-to-one version of this problem has already been studied by Dughmi and Ghosh [4] in a setting where the weights w_{ij} are public knowledge, and the agents only report the tasks they are able to perform. We study here the case where the weights are private data. We are interested in the best approximation ratios that can be achieved by a truthful mechanism. In particular, we investigate the problem under various assumptions on the way the agents can misreport the weights.

Key words: Algorithmic game theory, truthful mechanism without payment, approximation algorithm, many-to-many assignment problem.

1 Introduction

We study here many-to-many assignment problems, where a set of tasks $\{t_1, \dots, t_m\}$ are assigned to a set of agents $\{a_1, \dots, a_n\}$. Let $x_{ij} = 1$ if task t_j is assigned to agent a_i . We tackle the case where one assigns p tasks per agent and q agents per task, i.e. $\sum_j x_{ij} = p$ and $\sum_i x_{ij} = q$. A weight $w_{ij} \geq 0$ can represent the interest shown by agent a_i to perform task t_j , but also an attendance level, a bandwidth, etc. The aim of the optimization problem is to assign the tasks to the agents so as to maximize some social objective, e.g. $\sum_{i,j} w_{ij}x_{ij}$ (in an *utilitarian* setting) or $\min_i \sum_j w_{ij}x_{ij}$ (in an *egalitarian* setting). This type of problem can occur in various circumstances:

- *Assignment of papers to reviewers for a conference.* Consider a program committee chair who has to assign the m submitted papers (the tasks) to n reviewers (the agents). Each reviewer is assigned p papers and each paper

^{*} This research has been supported by the project ANR-09-JCJC-0066-01 “Combinatorial Optimization with Competing Agents” (COCA).

is reviewed by q reviewers. In order to perform the assignment, during the bidding phase, for each paper t_j each reviewer a_i indicates an interest score w_{ij} on the conference system.

- *Assignment of public resources to non-profit associations.* Consider a town council who allocates m time slots/room locations (the tasks) to n sports or cultural associations (the agents). There are q time slots in a week (m/q room locations) and each association has the possibility of being assigned p time slots/room locations per week (once for all). Furthermore, for each time slot t_j , each association a_i indicates the expected number w_{ij} of members that would attend the activity.
- *Assignment of time slots on an earth observation satellite.* Consider an earth observation satellite shared among n institutions (the agents). In the daily management of the satellite, there are m time slots (the tasks) in a day and a limited number q of possible accesses to the same slot. For a given institution, there is a maximum of p accesses per day and one access per time slot. Furthermore, for each time slot t_j , each institution a_i indicates the required bandwidth w_{ij} .

In all these circumstances, based on the bids, the tasks are assigned to the agents so as to maximize the social welfare (depending on the social objective). Once the bids are known, the optimization problems involved are solvable in polynomial time, except for the many-to-many case in the egalitarian setting (strongly NP-hard by reduction from the 3-partition problem, see Section 4). These optimization problems have also been investigated from the fairness viewpoint in at least two of the three situations described above, namely for the fair paper assignment problem (see e.g. [5, 6, 11]) and for the fair sharing of a common satellite resource (see e.g. [10]). Nevertheless, since the weights are private data, an agent might have an incentive to misreport them in order to increase her individual satisfaction (the individual satisfaction of an agent a_i is $\sum_j w_{ij}x_{ij}$).

The aim of the present article is precisely to study *truthful* mechanisms for many-to-many assignment problems. A mechanism is truthful if no agent i can benefit by misreporting her weights w_{ij} . At the same time, the mechanism should guarantee a reasonable approximation to the optimal social objective. Note that we focus on mechanisms without payment, since payments are not always possible or desirable [15] (as in most examples above).

Truthful mechanisms dedicated to the one-to-one version of the problem have already been studied by Dughmi and Ghosh [4] for the utilitarian setting in a restricted case where, for each pair (i, j) , it is public knowledge that the value of task t_j for agent a_i is either w_{ij} or 0, but agent a_i holds private which of those is the case for each t_j . The authors justify considering such a discrete valuation model by observing that, assuming the weights are private, no mechanism can outperform the trivial one which allocates the tasks uniformly at random, achieving an approximation ratio of n (as a consequence of classical results of Vickrey concerning single item auctions [16]). Nevertheless, we believe that some interesting results are still possible with private weights.

First, the result given by Dughmi and Ghosh [4] does not hold in the egalitarian setting (in this case, returning a random assignment yields an approximation ratio of $n!$). Second, for both the utilitarian and the egalitarian settings, positive results are possible if one assumes some restrictions on the way the agents misreport their preferences. We consider two types of restrictions: the case where the agents cannot overbid (i.e. they cannot bid weights that are larger than their true weights), and the case where the agents cannot underbid (i.e. they cannot bid weights that are smaller than their true weights). The former case occurs for instance in the assignment of public resources to non-profit associations, as described above: as soon as the actual attendance level at the various activities is controlled afterwards, an association cannot overbid, at the risk of losing credibility. The latter case occurs for instance in the assignment of time slots on an earth observation satellite, as described above: as soon as the bandwidth allocated to an institution is equal to the indicated weight, there is no interest for an institution to bid a value under the required bandwidth (an undersized bandwidth is equivalent to a bandwidth 0 for the institution). The rationale for these restrictions is based on the assumption that the agents underbid in order to avoid a task, and conversely they overbid in order to obtain a task. This assumption holds for all the mechanisms we propose. Note that the restriction on the way the agents misreport their preferences is related to the idea of so-called *mechanisms with verification* [13]. Mechanisms with verification have been studied both with payment [14, 9] and without payment [1, 8].

Our Contribution. In Section 2, we study the case where the weights on the edges are unrestricted. In particular, we provide an n -approximate randomized mechanism for the egalitarian setting and we show that there is no $(n/q - \varepsilon)$ -algorithm for this problem. In Section 3, we study the case where the agents cannot overbid. This assumption does not change anything for the egalitarian setting, but it enables to design a 2-approximate truthful mechanism for the utilitarian setting (and we show this is the best ratio that can be achieved by a truthful algorithm). In Section 4, we study the case where the agents cannot underbid. Conversely to the previous case, this assumption does not change anything for the utilitarian setting, but it enables to design an optimal truthful mechanism for the egalitarian setting. Our results are summarized in Table 1, where we indicate lower (LB) and upper (UB) bounds on the approximation ratio of a truthful mechanism. Notation “det” (resp. “rand”) stands for deterministic (resp. randomized) mechanisms. When no approximation guarantee can be achieved by a truthful mechanism, we write $LB = \infty$. Note that all positive results are based on polynomial-time algorithms, except the optimal truthful mechanism for agents that do not underbid in the egalitarian setting.

Preliminaries. As customary, we view the many-to-many assignment problem as a maximum weight b -matching problem in a complete bipartite graph G with vertex set $V = (A, T)$ (A for agents, T for tasks). Given degree constraints $b : V \rightarrow \mathbb{N}$ for the vertices, a b -matching is a set of edges M such that for all $v \in V$ the number of edges in M incident to v , denoted by $\deg_M(v)$, is at most $b(v)$. Every agent a_i has to perform (at most) $b(a_i) = p$ tasks, and every

	<i>utilitarian setting</i>	<i>egalitarian setting</i>
no restriction	det: $LB = \infty$ rand: $LB = (\frac{n}{q} - \varepsilon), UB = \frac{n}{q}$	det: $LB = \infty$ rand: $LB = (\frac{n}{q} - \varepsilon), UB = n$
no overbidding	det: $LB = 2 - \varepsilon$ [4], $UB = 2$ rand: $LB = (\frac{4}{3} - \varepsilon), UB = 2$	det: $LB = \infty$ rand: $LB = (\frac{n}{q} - \varepsilon), UB = n$
no underbidding	det: $LB = \infty$ rand: $LB = (\frac{n}{q} - \varepsilon), UB = \frac{n}{q}$	det: $UB = 1$

Table 1. An overview of our results.

task t_j has to be assigned to (at most) $b(t_j) = q$ agents. There are $n = |A|$ agents and $m = |T|$ tasks, and we assume that $np = mq$ (otherwise some agents would perform less than p tasks or some tasks would be assigned to less than q agents). The weight $w_{ij} \geq 0$ on edge $\{a_i, t_j\}$ represents the interest (value) of agent a_i for task t_j . Note that, since $np = mq$ and by positivity of the weights and completeness of the graph, constraints $\deg_M(v) \leq b(v) \forall v$ are equivalent to $\deg_M(v) = b(v) \forall v$ (if some constraints are not saturated, it is always possible to add (an) edge(s) to the b -matching to saturate them). Let $M(i)$ denote the set of indices of the p tasks assigned to agent a_i in the b -matching M . The aim of each agent a_i is to maximize $\sum_{j \in M(i)} w_{ij}$. A deterministic mechanism (algorithm) is truthful if no agent a_i has an incentive to misreport the weights w_{ij} ($j = 1, \dots, n$): $\sum_{j \in M(i)} w_{ij} \geq \sum_{j \in M'(i)} w_{ij}$ where M (resp. M') is the b -matching returned by the mechanism if agent a_i truthfully reports (resp. misreports) her weights (*ceteris paribus*). For a randomized mechanism, we require truthfulness in expectation: $\mathbb{E}[\sum_{j \in M(i)} w_{ij}] \geq \mathbb{E}[\sum_{j \in M'(i)} w_{ij}]$.

We recall that our aim is to obtain truthful mechanisms returning a b -matching M which maximizes one of the following social objective functions:

- $w(M) = \sum_{i=1}^n \sum_{j \in M(i)} w_{ij}$ (in the utilitarian setting),
- or $w(M) = \min_{i \in \{1, \dots, n\}} \sum_{j \in M(i)} w_{ij}$ (in the egalitarian setting).

A deterministic (resp. randomized) mechanism is said to be c -approximate if, for any instance, $w(M^*)/w(M) \leq c$ (resp. $w(M^*)/\mathbb{E}[w(M)] \leq c$) where M is the b -matching returned by the mechanism, M^* is an optimal b -matching and $\mathbb{E}[w(M)]$ is the expected weight of M .

2 If Agents Have No Restriction on the Way They Bid

In this section, we show that no performance guarantee can be achieved by a truthful deterministic mechanism, in both the utilitarian setting and the egalitarian setting. Furthermore, we show that there are n -approximate randomized truthful mechanisms (in both settings).

2.1 Utilitarian Setting

We show that no positive result can be expected for a deterministic mechanism.

Theorem 1 *In the utilitarian setting, for any $c > 1$, there is no c -approximate truthful deterministic mechanism, even if $n = m = 2$.*

Proof. Let $c > 1$, $n = m$ and $p = q = 1$ (the b -matching problem reduces to a matching problem). Assume that there is a c -approximate truthful mechanism and consider the graph with two agents a_1, a_2 and two tasks t_1, t_2 . Let $\gamma > c$ and $w_{11} = \gamma$, $w_{12} = 0$, $w_{21} = 1$ and $w_{22} = 0$. There are only two perfect matchings $M_1 = \{\{a_1, t_1\}, \{a_2, t_2\}\}$ and $M_2 = \{\{a_1, t_2\}, \{a_2, t_1\}\}$ whose weights are γ and 1 respectively. Since $\gamma > c$, the mechanism must return M_1 .

Now, take the situation where $w_{21} = \gamma^2$. The weights of M_1 and M_2 are then γ and γ^2 . Again, a c -approximate mechanism must return M_2 since otherwise $\frac{w(M_2)}{w(M_1)} = \frac{\gamma^2}{\gamma} > c$. Consequently, in the first situation, agent a_2 has incentive to declare a false weight $w_{21} = \gamma^2$ in order to get task t_1 , with value 1 instead of 0 (with the initial task t_2). ■

Note that as a consequence there is no $f(n, m)$ -approximate deterministic mechanism, for any function f .

In the randomized case, as mentioned earlier, Dughmi and Ghosh [4] have observed that no truthful mechanism can perform better than the mechanism returning a random assignment in the one-to-one case (whose approximation ratio is n since each edge of an optimal matching has a probability $1/n$ to belong to the returned matching). Their statement is based on a reference to a seminal paper by Vickrey [16], whose scope is much broader than the simple assignment problem. This result generalizes to the many-to-many case. First, note that the approximation ratio of the mechanism returning a random b -matching is $n/q = m/p$ since each edge of an optimal b -matching has a probability $q/n = p/m$ to belong to the returned b -matching. The following theorem shows that this is the best we can do:

Theorem 2 *In the utilitarian setting, for any n , any q , any $\varepsilon > 0$, there is no $(\frac{n}{q} - \varepsilon)$ -approximate truthful randomized mechanism.*

Proof. Let γ such that $1/\gamma < 1/(n/q - \varepsilon) - q/n$. Consider the graph with n agents and n tasks where $w_{11} = \gamma + \frac{1}{q} - 1$, $w_{i1} = 1/q$ for $i \geq 2$, and all the other weights are 0 (note that $p = q$ since $n = m$). Note that the optimal solution has value $w_{11} + (q-1) \times 1/q = \gamma$. Let p_i be the probability that edge (a_i, t_1) is taken in the matching returned by the randomized mechanism. Then since any b -matching without edge (a_1, t_1) (resp. with edge (a_1, t_1)) has value at most 1 (resp. γ), we get that the expected value of the solution is at most $p_1\gamma + (1 - p_1) \leq p_1\gamma + 1$. To achieve ratio $n/q - \varepsilon$, we need $p_1\gamma + (1 - p_1) \geq \gamma/(n/q - \varepsilon)$ and thus $p_1\gamma + 1 \geq \gamma/(n/q - \varepsilon)$. Hence $p_1 \geq 1/(n/q - \varepsilon) - 1/\gamma > q/n$ (by the choice of γ). Since $\sum p_i = q$ (t_1 is adjacent to q edges in the returned b -matching), there exists i such that $p_i < q/n$, say $i = 2$ wlog (note that $i \neq 1$ since $p_1 > q/n$).

Now, consider the same situation as before but with $w_{21} = \gamma^2$. By using the same argument as above, we see that the mechanism has to choose edge (a_2, t_1) with probability greater than q/n .

Then, in the first situation, if agent 2 truthfully reports her weight w_{21} she gets t_1 with probability $p_2 < q/n$, while if she reports a weight γ^2 for this edge she gets t_1 with probability greater than q/n . The mechanism is not truthful. ■

2.2 Egalitarian Setting

Similarly to the utilitarian setting, we show that no positive result can be expected for a deterministic mechanism.

Theorem 3 *In the egalitarian setting, for any $c > 1$, there is no c -approximate truthful deterministic mechanism, even if $n = 2$, $m = 2$ and $p = q = 1$.*

Proof. Let $c > 1$, $n = m = 2$ and $p = q = 1$. By contradiction, assume that there is a c -approximate truthful deterministic mechanism and consider the graph with two agents a_1, a_2 and two tasks t_1, t_2 . Let $\gamma > c$ and $w_{11} = \gamma + 1$, $w_{12} = \gamma$, $w_{21} = \gamma + 1$ and $w_{22} = \gamma$. There are only two perfect matchings $M_1 = \{\{a_1, t_1\}, \{a_2, t_2\}\}$ and $M_2 = \{\{a_1, t_2\}, \{a_2, t_1\}\}$. Wlog, assume that the mechanism returns matching M_1 .

Now, consider the same situation as before, but $w_{22} = 1$. The weight of matching M_1 (resp. M_2) is 1 (resp. γ), where the weight of M_1 (resp. M_2) is the minimal weight of an edge in M_1 (resp. M_2). Since the mechanism is c -approximate, it must return matching M_2 (because $\gamma > c$).

Then, in the first situation, agent a_2 has incentive to bid a false weight $w_{22} = 1$ in order to get task t_1 , with value $\gamma + 1$ instead of γ (with the initial task t_2). Therefore, the mechanism is not truthful. ■

Note that as a consequence there is no $f(n, m)$ -approximate truthful deterministic mechanism, for any function f .

In the randomized case, note that, in contrast to the utilitarian setting, returning a random assignment is not n/q -approximate anymore for the egalitarian setting (consider an instance where $p = q = 1$, and thus $n = m$, where all the weights are equal to 0, except weights $w_{ii} = 1$ for $i = 1, \dots, n$: there is only one matching over $n!$ that has weight 1, all the others having weight 0). However, we now show that it is possible to get an n -approximation thanks to the following truthful randomized mechanism:

Let M^* denote an optimal b -matching for the egalitarian setting. Return with probability $1/n$, for each $k \in \{0, \dots, n-1\}$, the b -matching where tasks with indices in $M^*((i+k) \bmod n)$ are assigned to agent a_i ($i = 1, \dots, n$).

Theorem 4 *In the egalitarian setting, the above mechanism is truthful and n -approximate.*

Proof. This mechanism is truthful since, whatever the bids, every agent a_i has the same expected value $\sum_{k=1}^n \frac{1}{n} \sum_{j \in M^*(k)} w_{ij} = \frac{q}{n} \sum_j w_{ij}$ (since each task is assigned to q agents in M^*). For $k = 0$, the mechanism returns M^* . The b -matching M^* is thus returned with probability $1/n$, and the mechanism is n -approximate. ■

Note that this mechanism is also a truthful n -approximate mechanism for the utilitarian setting. The following theorem shows that no truthful mechanism can have an approximation ratio better than n/q (the previous truthful mechanism achieves therefore the best possible ratio when $q = 1$):

Theorem 5 *In the egalitarian setting, for any n , any q and any $\varepsilon > 0$, there is no $(\frac{n}{q} - \varepsilon)$ -approximate truthful randomized mechanism, even if $p = q$.*

Proof. Let us consider a truthful randomized mechanism. Consider the graph with n agents and n tasks ($p = q$) where $w_{i1} = \gamma \geq 1$ for $i = 1, \dots, n$, and all the other weights are $1/q$ (see Figure 1, Situation 1). Since t_1 is assigned to q agents, there exists an agent, say a_1 wlog., such that the randomized mechanism returns a b -matching containing edge $\{a_1, t_1\}$ with probability smaller than or equal to q/n .




Fig. 1. Illustration of Theorem 5. The sharp edges have weight γ , the dotted edges have weight 0, and the unrepresented edges have weight $1/q$.

Now, consider the same situation as before but $w_{11} = \gamma$ and $w_{1j} = 0$ for $j = 2, \dots, n$ (see Figure 1, Situation 2). Let p_1 be the probability that the randomized mechanism returns a b -matching containing edge $\{a_1, t_1\}$ for this instance. The expected weight of the returned b -matching is $p_1 \times 1 + (1 - p_1) \times 0 = p_1$ (the only non-zero b -matchings are the ones where edges $\{a_1, t_1\}$ is chosen).

Since the mechanism is truthful, $p_1 \leq q/n$ (otherwise, in situation 1, agent 1 would have incentive to bid false values in order to be in situation 2). Thus, in situation 2, since $p_1 \leq q/n$, the expected weight of the returned b -matching is at most q/n , while the optimal matching has weight 1. The mechanism is therefore at most n/q -approximate. ■

3 If Agents Do Not Overbid

In this section, we assume that the agents cannot bid weights that are strictly larger than their true weights. This assumption does not change anything in the egalitarian setting, since the agents do not overbid in the situation used to

establish the lower bound in the unrestricted case. However, we can provide a 2-approximate deterministic truthful mechanism for the utilitarian setting.

Sort the edges by non-increasing weights. Let $M = \emptyset$ and (e_1, \dots, e_m) denote the sorted list of the edges. For i from 1 to m , if $M \cup \{e_i\}$ is a b -matching then $M = M \cup \{e_i\}$. Return M .

This greedy b -matching algorithm has been introduced by Avis [2] for the case $p = q = 1$, who has shown that it is 2-approximate. Mestre [12] has shown that this approximation ratio also holds for b -matchings.

Note that the tie-breaking rule (to decide an ordering over edges of equal weight) matters for the truthfulness of the mechanism. A convenient and simple tie-breaking rule is the following one: if $\{a_i, t_j\}$ and $\{a_{i'}, t_{j'}\}$ are such that $w_{ij} = w_{i'j'}$, $\{a_i, t_j\}$ is ranked before $\{a_{i'}, t_{j'}\}$ if $i < i'$ or if $i = i'$ and $j < j'$.

Now, we are able to show that the greedy b -matching algorithm is truthful.

Theorem 6 *In the utilitarian setting, if the agents cannot overbid, the greedy b -matching algorithm is a truthful 2-approximate mechanism.*

Proof. We have already said that it returns a 2-approximate matching. We now show that it is also truthful.

By contradiction, assume that agent a_i has incentive to lie on k weights. Let M (resp., M') be the b -matching returned by the algorithm with weights w (resp. w'), i.e., when agent a_i does not (resp. does) misreport her weights. By abuse of notation, we denote by $M(i)$ (resp. $M'(i)$) the set of edges incident to a_i in M (resp. M'). Let $M(i) \setminus M'(i) = \{\{a_i, t_{\pi(1)}\}, \dots, \{a_i, t_{\pi(k)}\}\}$, where $\{a_i, t_{\pi(j)}\}$ is the j^{th} edge of $M(i) \setminus M'(i)$ examined by the algorithm for weights w .

One proceeds as follows to decompose $M \Delta M' = (M \setminus M') \cup (M' \setminus M)$ into k edge-disjoint cycles C_1, \dots, C_k alternating one edge in $M \setminus M'$ and one edge in $M' \setminus M$, each cycle C_j including edge $\{a_i, t_{\pi(j)}\}$: initially, $C_1 = \emptyset$; note that $\{a_i, t_{\pi(1)}\}$ is necessarily the first edge in $M \Delta M'$ examined by the algorithm for weights w (provided the agents cannot overbid); insert $\{a_i, t_{\pi(1)}\}$ in C_1 ; then extend C_1 with the first edge in $M'(t_{\pi(1)}) \setminus M(t_{\pi(1)})$ examined by the algorithm; this latter edge concerns a given agent $a_{i'}$; extend C_1 with the first edge in $M(i') \setminus M'(i')$; and so on until a cycle is created. This is cycle C_1 . One iterates to obtain the other cycles C_2, \dots, C_k .

Let $\{a_i, t_{\mu(j)}\}$ denote the single edge in $C_j \cap M'(i)$. Since $\{a_i, t_{\mu(j)}\}$ is examined after $\{a_i, t_{\pi(j)}\}$ in cycle C_j , one has $w_{i\mu(j)} \leq w_{i\pi(j)}$. By summing up these inequalities one obtains $\sum_{j \in M'(i)} w_{ij} \leq \sum_{j \in M(i)} w_{ij}$. Thus agent a_i has no incentive to underbid, and the mechanism is truthful. ■

The following theorem, due to Dughmi and Ghosh in the simpler case where $p = q = 1$, shows that no truthful deterministic mechanism can have a better approximation ratio:

Theorem 7 ([4]) *In the utilitarian setting, for any $\varepsilon \in (0, 1)$, there is no $(2-\varepsilon)$ -approximate truthful deterministic mechanism for the utilitarian setting, even if $n = 2$ and $p = q = 1$.*

A similar proof makes possible to establish the following negative result for randomized mechanisms:

Theorem 8 *In the utilitarian setting, for any $\varepsilon \in (0, 1/3)$, there is no $(4/3 - \varepsilon)$ -approximate truthful randomized mechanism for the utilitarian setting, even if $n = 2$ and $p = q = 1$.*

Proof. Let $\varepsilon \in (0, 1/3)$. Consider a $\gamma > 0$ such that $\frac{2+\gamma}{3/2+\gamma} > \frac{4}{3} - \varepsilon$.

By contradiction, suppose that there is a $(4/3 - \varepsilon)$ -approximate truthful mechanism and consider the graph with 2 agents a_1, a_2 and 2 tasks t_1, t_2 . Let $w_{11} = w_{21} = 1 + \gamma$ and $w_{12} = w_{22} = 1$. There are only two perfect matchings $M_1 = \{\{a_1, t_1\}, \{a_2, t_2\}\}$ and $M_2 = \{\{a_1, t_2\}, \{a_2, t_1\}\}$ with same weights and by symmetry, wlog., consider that the mechanism returns the matching M_1 with a probability $p_1 \geq 1/2$ (see Figure 2 case (a)). In solid (resp., dotted) lines is the matching M_1 (resp., M_2).

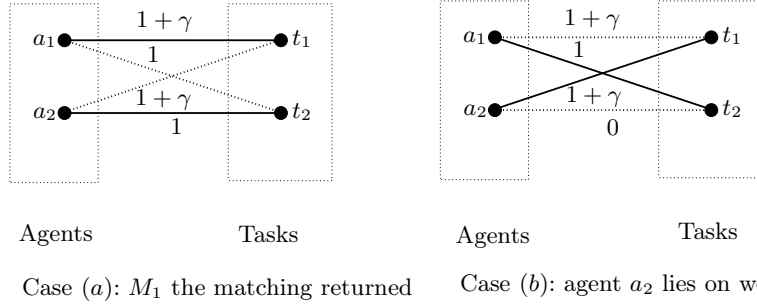


Fig. 2. Illustration of Theorem 8.

Now, take the situation where $w_{22} = 0$ (see Figure 2, case (b)). The weights of matching M_1 (resp. M_2) is $1 + \gamma$ (resp. $2 + \gamma$). By truthfulness, the mechanism should return matching M_1 with probability at least p_1 (otherwise in the initial situation agent 2 would have incentive to bid a false value for edge $\{a_2, t_2\}$). The expected weight of the matching returned by the mechanism in situation 2 is thus $p_1(1 + \gamma) + (1 - p_1)(2 + \gamma) = 2 + \gamma - p_1$. Since $p_1 \geq 1/2$, this weight is smaller than or equal to $3/2 + \gamma$. Since the weight of the optimal matching is $2 + \gamma$, the expected approximation ratio of the mechanism is larger then or equal to $\frac{w(M^*)}{w(M_1)} = \frac{2+\gamma}{3/2+\gamma} > \frac{4}{3} - \varepsilon$. Thus the mechanism is not $(4/3 - \varepsilon)$ -approximate. ■

4 If Agents Do Not Underbid

In this section, we assume that the agents cannot bid weights that are strictly smaller than their true weights. This assumption does not change anything in the utilitarian setting, since the agents do not underbid in the situations used to establish the lower bound in the unrestricted case.

Let us then consider the egalitarian setting. We will show that there is an optimal truthful mechanism. As a first remark, the optimization problem of

finding an optimal b -matching is strongly **NP**-hard: take the 3-partition problem where given a set of $3n$ positive integers $\{s_1, \dots, s_{3n}\}$ of total sum nB we ask whether there exist n subsets of weight B , each of them consisting of exactly three elements. Build the graph with n agents and $3n$ tasks (one task t_i for each integer s_i), the edges adjacent to a task t_i being weighted by s_i . It is easy to see that the answer to the 3-partition problem is yes if and only if there exists a b -matching of (egalitarian) value B .

So an optimal truthful mechanism is not polynomial time (unless **P=NP**), even if $p = 3$. We consider the following mechanism, where P is an upper bound on the optimal value (say P equals the sum of the p maximum edge weights adjacent to agent a_1 for instance).

Set $W \leftarrow P$. While no solution has been found:

- Consider for each agent a_i the set $S_i(W)$ of all sets T' of p tasks such that $\sum_{t_j \in T'} w_{ij} \geq W$.
 - Use an algorithm **A** to determine whether there exists a way to assign to each a_i a set of tasks in $S_i(W)$ such that each task is assigned to q agents.
- If a solution is found, output it. Otherwise, set $W \leftarrow W - 1$.

Note that the mechanism terminates: there necessarily exists a feasible assignment for $W = 0$ since the graph is then complete.

Stated like this, the mechanism might not be truthful. More precisely, to make it truthful we have to be more careful on the way sets in $S_i(W)$ are assigned to agents. Informally, by overbidding an agent may add new sets of tasks in $S_i(W)$. We have to be sure that by doing this she will not be better off. To ensure this, we shall use an algorithm **A** which has the following *stability* property:

Stability: if **A** assigns on an instance I (defined by the set of agents and the set $S_j(W)$ for each agent a_j) the set of tasks $M(i) \in S_i(W)$ to agent a_i , then on an instance I' where $S'_i(W) \supseteq S_i(W)$ (and $S'_j(W) = S_j(W)$ for each $j \neq i$), **A** assigns to a_i either $M(i)$ or a set in $S'_i(W) \setminus S_i(W)$.

It is not difficult to see that there exists a stable algorithm **A**. Indeed, just compute (if any) an assignment of maximal total weight, where the weights of sets in $S_1(W), \dots, S_n(W)$ are such that any two assignments have different weights, these weights depending only on n (not on W)³. This way, when sets are added to $S_i(W)$, if the (unique) maximum weight assignment has not changed a_i receives $M(i)$, otherwise the new maximum weight assignment uses a set in $S'_i(W) \setminus S_i(W)$ for a_i , thus ensuring stability.

Theorem 9 *In the egalitarian setting, if the agents cannot underbid, the mechanism with a stable algorithm **A** is both optimal and truthful.*

³ These weights can be defined as follows: the weight of a set $M(k)$ in $S_k(W)$ of tasks is $2^{\alpha_{M(k)}}$ where $\alpha_{M(k)} \in \mathbb{N}$ is different for any two sets of tasks (for instance give a weight $v_{ij} = 2^{i(n-1)+(j-1)}$ to each edge (i, j) in G and define $\alpha_{M(k)}$ as the total weight $\sum_{j \in M(k)} v_{kj}$).

Proof. The b -matching returned by the algorithm at the iteration W has value at least W (by construction). Since we consider W by decreasing value, the mechanism is clearly optimal. For truthfulness, consider a graph G with weights w . Suppose that the mechanism has found a b -matching M when the threshold is W , assigning the set of tasks $M(i)$ to agent a_i . We have $\sum_{j \in M(i)} w_{ij} \geq W$. Suppose that agent a_i reports weights $w'_{ij}, j = 1, \dots, n$ where $w'_{ij} \geq w_{ij}$. On this new instance G weighted by w' (where $w'_{kj} = w_{kj}$ for $k \neq i$), suppose that the mechanism has found a b -matching M' for threshold W' . Since no underbid is allowed, $W' \geq W$. Hence, two cases may occur.

- If $W' > W$, then necessarily in M' agent a_i receives a set $M'(i)$ where $\sum_{j \in M'(i)} w_{ij} \leq W$. Otherwise, a b -matching would have been found in G weighted by w with threshold (at least) $W + 1$.

- If $W' = W$, then the set $S'_i(W)$ (of the possible sets of p tasks for a_i according to weights w') is obtained from $S_i(W)$ (with weights w) by adding all the subsets of weight less than W according to w , but at least W according to w' . Hence, by the stability property, in M' agent a_i either receives $M(i)$, or one of the added sets whose weight (according to w) is smaller than W . In both cases, $\sum_{j \in M'(i)} w_{ij} \leq \sum_{j \in M(i)} w_{ij}$.

In both cases, agent a_i is not better off, and the mechanism is truthful. ■

In the particular case $p = q = 1$, the problem consists in finding an egalitarian matching and is polynomial time. Our mechanism reduces indeed to finding at each step W if there exists a perfect matching (polynomial time problem) in the graph consisting of edges of weight at least W (using a stable perfect matching algorithm); this is the threshold method [7], which is thus truthful by Theorem 9.

5 Final Remarks

Concerning the unrestricted case, note that the simple truthful mechanism that consists in considering the agents in a random order, and assigning to each of them her p preferred tasks among the available ones (random round robin), is n -approximate in the utilitarian setting⁴. Despite the fact that its performance guarantee is the same as the mechanism returning a random assignment, the random round robin mechanism should be preferred since it is likely to return much better b -matchings in practice.

In order to strengthen the result concerning the case where the agents do not underbid, it would be interesting to investigate what approximation can be achieved truthfully in polynomial time. Besides, another natural research direction is to investigate the group strategyproofness of the mechanisms presented in the paper: a mechanism is *group strategyproof* if for every group of

⁴ Let W_i be the total weight of the p edges of maximum weight which are adjacent to a_i . The expected cost of the solution returned is at least $\sum_{i=1}^n (1/n)W_i$, and the mechanism is clearly n -approximate. The bound is tight when all the edges have weight 0, except the edges adjacent to the same p tasks which have weight $M \gg 1$ for the edges adjacent to a_1 , and weight 1 for the edges adjacent to the other agents.

agents $A' \subseteq A$ and every weights w' such that $w'_{ij} = w_{ij}$ if $i \notin A'$, one has $\sum_{j \in M'(i)} w_{ij} \leq \sum_{j \in M(i)} w_{ij}$, where M (resp. M') is the b -matching returned by the mechanism for weights w (resp. w'). A third research direction is to study other restrictions regarding the possible ways that the agents can lie, while still showing positive approximation guarantees. For instance, what happens if $\sum_j w_{ij}$ is a constant for all the agents? (Think of a conference system that gives the same capital of points to each reviewer) Another possible domain restriction is the following: when attributing time slots to agents, one can assume that the preferences of the agent are single-peaked [3], i.e. each agent has an ideal time slot and her preferences decrease when moving away from the ideal point. It would be interesting to study dedicated truthful mechanisms taking advantage of these domain restrictions.

Acknowledgments. We wish to thank the anonymous reviewers for their fruitful comments on an earlier version of the paper.

References

1. E. Angel, E. Bampis, and F. Pascual. Truthful algorithms for scheduling selfish tasks on parallel machines. *Theoretical Computer Science*, 369(1-3):157–168, 2006.
2. D. Avis. A survey of heuristics for the weighted matching problem. *Networks*, 13:475–493, 1983.
3. D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.
4. S. Dughmi and A. Ghosh. Truthful assignment without money. In *ACM Conference on Electronic Commerce*, pages 325–334, 2010.
5. N. Garg, T. Kavitha, A. Kumar, K. Mehlhorn, and J. Mestre. Assigning papers to referees. *Algorithmica*, 58(1):119–136, 2010.
6. B. Golden and P. Perny. Infinite order Lorenz dominance for fair multiagent optimization. In *AAMAS*, pages 383–390, 2010.
7. O. Gross. The bottleneck assignment problem. Technical Report P-1630, The Rand Corporation, Sta. Monica, CA, 1959.
8. Elias Koutsoupias. Scheduling without payments. In *SAGT*, pages 143–153, 2011.
9. P. Krysta and C. Ventre. Combinatorial auctions with verification are tractable. In *ESA (2)*, pages 39–50, 2010.
10. M. Lemaître, G. Verfaillie, H. Fargier, J. Lang, N. Bataille, and J.M. Lachiver. Equitable allocation of earth observing satellites resources. In *In Proc of 5th ONERA-DLR Aerospace Symposium (ODAS'03)*, 2003.
11. J. Lesca and P. Perny. LP solvable models for multiagent fair allocation problems. In *ECAI*, pages 393–398, 2010.
12. J. Mestre. Greedy in approximation algorithms. In *ESA*, pages 528–539, 2006.
13. N. Nisan and A. Ronen. Algorithmic mechanism design (extended abstract). In *STOC*, pages 129–140, 1999.
14. P. Penna and C. Ventre. Optimal collusion-resistant mechanisms with verification. In *ACM Conference on Electronic Commerce*, pages 147–156, 2009.
15. J. Schummer and R.V. Vohra. Mechanism Design without Money. In N. Nisan, T. Roughgarden, E. Tardos, and V.V. Vazirani, editors, *Algorithmic Game Theory*, pages 243–266. Cambridge University Press, 2007.
16. W. Vickrey. Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance*, pages 8–37, 1961.