Facilitating Matches on Allocation Platforms

Yohai Trabelsi [a,*,1], **Abhijin Adiga**[b], **Yonatan Aumann**[c], **Sarit Kraus**[c] **and S. S. Ravi**[b]

[a]Harvard University
[b]Biocomplexity Institute, University of Virginia
[c]Dept. of Computer Science, Bar-Ilan University

Abstract.
We consider a setting where goods are allocated to agents by way of an allocation platform (e.g., a matching platform). An "allocation facilitator" aims to increase the overall utility/social-good of the allocation by encouraging (some of the) agents to relax (some of) their restrictions. At the same time, the advice must not hurt agents who would otherwise be better off. Additionally, the facilitator may be constrained by a "bound" (a.k.a. 'budget'), limiting the number and/or type of restrictions it may seek to relax. We consider the facilitator's optimization problem of choosing an optimal set of restrictions to request to relax under the aforementioned constraints. Our contributions are three-fold: (i) We provide a formal definition of the problem, including the participation guarantees to which the facilitator should adhere. We define a hierarchy of participation guarantees and also consider several social-good functions. (ii) We provide polynomial algorithms for solving various versions of the associated optimization problems, including one-to-one and many-to-one allocation settings. (iii) We demonstrate the benefits of such facilitation and relaxation, and the implications of the different participation guarantees, using extensive experimentation on three real-world datasets.

1 Introduction

Recently, allocation-platforms/matching-platforms, which allocate resources of one sort or another to users, are being deployed for a variety of applications in both the public and private sectors, including in welfare and social services [17, 23, 3]. Some examples are allocating home healthcare demand with service providers [20], on-demand housekeeping platforms [34], government platforms for providing housing assistance to homeless individuals [25], ride-sharing platforms [4], sharing parking spaces [28], and volunteer matching platforms [27]. In these platforms, users specify their resource requirements and constraints, and the platforms aim to optimize the allocation of resources to users where resources can be shared among several users.

The key stakeholders in the resultant allocation are clearly the users. In many cases, however, there may be additional stakeholders. For example, local welfare authorities are justifiably interested in increasing the number of homeless individuals awarded housing assistance (see e.g., [7]), and university administration is interested in maximizing the number of courses for which classrooms have been successfully allocated (see e.g., [14]). At times, these stakeholders may be able to directly determine, or make changes to the allocation, but more often than not, the allocation procedure itself is fixed - for regulatory, commercial, or technical reasons (see e.g., [30]). For example, New York City offers affordable housing opportunities through the *Housing Connect portal* [2], which serves as an allocation and matching platform. The algorithm used for this allocation is governed by multiple laws and regulations, including: (i) the federal Fair Housing Act, (ii) the NYC Human Rights Law, (iii)NYC HPD regulations, and (iv) NYC HDC rules. These regulations determine the allocation, which does not change. At the same time, the NYC Department of Housing Preservation and Development (HPD) established an advisory initiative, called the *Housing Ambassadors Program* [1], to help people navigate and use the allocation platform effectively. In their website, they emphasize that the Housing Ambassadors do not provide housing directly, and they cannot guarantee that an applicant will receive an affordable unit through the lottery.

In such cases, interested parties – which from now on we term *facilitators* – can still shape the resulting allocation by assisting and advising users in selecting the priorities and constraints they enter into the allocation platform. It is important to stress that such advice need not be viewed as a form of manipulation, neither of the platform nor of the users. Indeed, users frequently do not know how best to express their true constraints, and such interventions – if done right – can benefit all [26]. As such, we only consider *impartial* facilitators whose priorities are aligned with the overall social good, without any preference for one user or another. In this paper, we study advice provisioning with such impartial facilitators: what advice should they provide? What guarantees should/can be offered to the users, both those following the advice and those who do not?

In this paper, we assume that with appropriate guarantees, agents are willing to accept the advice. Providing them knowledge about whether they have been, or are currently, guaranteed a resource— either through the facilitator to ease constraints or the allocator as an act of transparency— further enhances their receptiveness.

For concreteness, we consider the following stylized model (see Figure 1 for an example). There is an allocation platform that, given (i) a set X of agents - each agent x_i with *demand level* d_i for the number of resources it needs, (ii) a set Y of resources, and (iii) a binary compatibility relation E between agents and resources, outputs a maximum allocation of resources to agents (where the maximum is in the sense of resource utilization). The compatibility relation E represents the compatibility *as provided to the platform by the agents*. The facilitator can advise agents to add additional compatibilities $\hat{E}$, thereby enlarging the input compatibility relation to $E^+ = E \cup \hat{E}$. If offered no advice, then $E^+ = E$. Adding any such new com-

* Corresponding Author. Email: yohai.trabelsi@gmail.com.
[1] Work done while the author was affiliated with Bar-Ilan University, Israel.

patibility $e \in \hat{E}$ is associated with some *discomfort level* $\rho(e)$. We consider a class of problems where the facilitator's objective is to maximize resource allocation with the constraint that the *aggregate cost* does not exceed a specified bound. Here, the aggregate cost is some function of all $\rho(e)$s, such as sum of all $\rho(e)$s or the count of strictly positive $\rho(e)$ values (which is the number of proposed relaxations). Figure 1 shows solutions for different functions. The incompatibility between a course and a classroom in Figure 1 can arise due to various factors such as seating capacity, commute distance, and accessibility issues. Optimal solutions for various scenarios (bound type and value—participation guarantee) are shown, all of which can be achieved by our framework. Expansions of the acronyms used in these scenarios are given under "Summary of Contributions". Formal definitions of the guarantees are provided in Section 3. Note that even for such a simple example, the solutions can be very different for different scenarios.

We propose the following two requirements to ensure that the proposed relaxation does not harm the agents while encouraging cooperation among them:

(i) *No agent is harmed:* Any agent that was guaranteed to be granted an allocation prior to the facilitator's actions, is also guaranteed so following the facilitator's actions.
(ii) *Participating agents benefit:* Any agent that is asked by the facilitator to add a compatibility (and does so) is guaranteed to be granted an allocation.

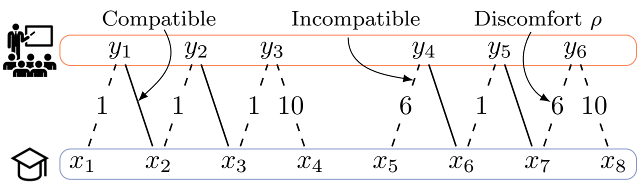

Bound	Size: 3	Total Cost: 15
Participation Guarantee	SNH-SB, WNH-WB, SNH-WB	WNH-WB
Facilitator's Advice	x_4–y_3, x_8–y_6	x_1–y_1, x_2–y_2, x_3–y_3, x_8–y_6

Figure 1. An example of a course–classroom allocation platform where every course needs to be matched to one classroom. Solid edges represent a compatible pair of a course and a classroom, dashed edges indicate an incompatible pair with a finite relaxation cost, and absence of an edge indicates an infinite relaxation cost incompatibility. An example that distinguishes between the different guarantees is in Figure 2.

It should be noted that the facilitator cannot secure these guarantees by directly determining the allocation, as the platform is not under its control. Rather, the facilitator must secure these guarantees by properly designing its advice.

Given such a setting, the facilitator's objective is to devise a set of incompatibilities that it will request agents to relax, which (i) maximizes the resultant allocation, while (ii) maintaining the overall cost within the given bound, and (iii) maintaining the no-harm and the participating agents guarantees. Agents are incentivized to follow the facilitator's recommendation as this guarantees them an allocation.

Summary of Contributions.
(a) Problem formulation: We formally define the problem, including

- Participation guarantees: We consider two forms of "no harm" participation guarantees: *strong no harm guarantee (SNH)*, and *weak no harm guarantee (WNH)*, and two forms of benefit to relaxers guarantees: *strong benefit guarantee (SB)*, and *weak benefit guarantee (WB)*. In both cases, the strong guarantee holds regardless of the number of agents adhering to the facilitator's advice, while the weak guarantee holds if all agents comply with the facilitator's recommendations. Using the defined guarantees, we define three guarantee combinations: SNH-SB, WNH-WB, and SNH-WB. The WNH-SB combination is omitted because, in practice, ensuring a strong guarantee for no harm is more important than for the benefit to relaxers.
- Aggregated cost: Inspired by Faliszewski and Rothe [13], we consider two functions for aggregating discomfort from relaxing individual incompatibilities: size (total relaxations) and total cost (sum of discomforts).
- We consider three allocation settings: (i) one-to-one where a resource is allocated to at most one agent and an agent is assigned at most one resource, (ii) many-to-one where a resource is allocated to at most one agent but an agent could be assigned multiple resources, and (iii) one-to-many where an agent is assigned at most one resource but a resource could be allocated to multiple agents (e.g., by sharing a classroom).

(b) Polynomial algorithms: We provide polynomial algorithms to solve each of the problem variants, for the three participation guarantees combinations, and the two aggregation functions. For the reason mentioned above, we omit the combination of weak no harm and strong benefit.
(c) Experimental study: We applied the devised algorithms to three real-world datasets, and conducted experiments in all three allocations settings (one-to-one, one-to-many and many-to-one). In each, we study the improvement in allocation sizes obtained by the facilitator under the different participation guarantees. We show that for all guarantees, a significant increase in allocation size is obtained. Comparing the performance of the different participation guarantees, we show that if all agents comply with the facilitator's advice, then the stronger guarantees result in somewhat smaller allocations than the weaker guarantees, but using the stronger guarantees is more robust to agents' failure to comply with the advice (as one would expect in practice).

2 Related Work

Resource allocation in multi-agent systems. Many references discuss how agents express their requirements, identify efficiently solvable allocation problems and provide methods for evaluating the corresponding algorithms (see e.g., [9, 16, 11]). Nguyen et al. [22] provide a good survey on the complexity and approximability of problems in this area. Zahedi et al. [36, 35] present a method where dissatisfied agents can challenge the proposed allocation using counterfactual queries. Relaxing the criteria for compatibility in kidney matching is studied in [24, 19]. However, their focus is on general criteria, not on specific agents (patients). Methods for active advice generation for a single agent appear in Trabelsi et al. [30]. A multi-round setting where several dissatisfied agents are given advice is presented in Trabelsi et al. [31]. While these papers focus on agent satisfaction, our work additionally emphasizes the role of the facilitator.

Participation in all maximum matchings. Costa [10] presented an algorithm to partition a graph's edge set into three subsets: edges that participate in all maximum matchings, in some maximum matchings, and in none of the maximum matchings. Irving et al. [18] showed how to efficiently compute the Dulmage–Mendelsohn decomposition [12] of a bipartite graph, enabling direct identification of the set of nodes participating in all maximum matchings. Zhang et al. [37] showed another algorithm for this task. However, these works do not consider how changing the graph edges affects this set.

Modifying the graph for improving the allocation. Boehmer et al. [6] considered bribery and external manipulations for providing participation guarantees to an agent pair in stable matchings. Chen and Csáji [8], Gokhale et al. [15] and Bobbio et al. [5] considered adjusting the resource capacities for having a many-to-one matching with some desired properties (e.g., a stable matching). However, these works do not consider the criteria of participation in all possible maximum allocations. Participation in all maximum allocations is preferred in our context, as the allocator can choose any maximum allocation and is not restricted to any other property (like stability).

3 Definitions and Problem Formulation

Here, we provide definitions for the one-to-one setting. The extension to the many-to-one case is provided in Section 5.

3.1 Preliminaries

The setting. The setting consists of sets X of agents, Y of resources, and $E \subseteq X \times Y$ of *compatible pairs*. Here, $(x, y) \in E$ means that agent x is willing to be allocated the resource y (without relaxing her preferences). Technically, the triple $G = (X, Y, E)$ is simply a bipartite graph.

Discomfort. For incompatible pairs $(x, y) \notin E$, there is a *discomfort function* $\rho : \overline{E} \to \mathbb{R}^+ \cup \{\infty\}$, where $\rho(x, y)$ reflects the "discomfort" for agent x of relaxing the incompatibility (x, y), that is, the discomfort that would be experienced by x if allocated to y. If $\rho(x, y) = \infty$ then y cannot be allocated to x, and (x, y) are deemed *totally incompatible*. We let $E_R = \{(x, y)|(x, y) \notin E, \rho(x, y) < \infty\}$ denote the set of *relaxable* incompatibilities.

Relaxations and cost bound. Given (X, Y, E) and ρ, we seek to relax some of the incompatibilities in order to increase the size of the resulting maximal allocation (produced by the allocation platform). Totally incompatible pairs cannot be relaxed. Thus, technically, a relaxation is a set $\hat{F} \subseteq E_R$. We denote by $X(\hat{F})$ the set of agents that participate in the relaxation $\hat{F}$. The *aggregate cost* of this relaxation is obtained by aggregating discomfort induced by $\hat{F}$. We consider two *aggregation functions*:

- Total Cost: $T(\hat{F}) := \sum_{(x,y)\in\hat{F}} \rho(x, y)$.
- Size: $S(\hat{F}) := |\hat{F}|$.

We assume that the facilitator has a *cost bound* β on the aggregate cost of the relaxation it chooses.

Allocations. In the one-to-one case, an allocation $M \subseteq F$ is a matching and a maximum allocation is a maximum cardinality matching in the graph $G = (X, Y, F)$. An allocation in the many-to-many, many-to-one, and one-to-many cases are defined similarly. Given a set $F \subseteq E$, we denote by $\mu(F)$ the size of a maximum allocation of F, and by $\Gamma(F) \subseteq X$ the set of agents that participate in all maximum allocations of F. If several such allocations exist, one of the solutions is picked arbitrarily.

Minimal relaxation. A relaxation $\hat{F}$ is minimal if $\mu(E \cup (\hat{F} \setminus \{e\})) < \mu(E \cup \hat{F})$, for any $e \in \hat{F}$.

3.2 Participation Guarantees

As explained in the introduction, the relaxation advice provided by the facilitator must provide guarantees both to the agents participating in the relaxation and to those not. The followings four guarantees are used for defining three guarantee combinations:

Strong No Harm: Any agent that participates in all maximum matchings prior to any relaxation continues to have this benefit after the relaxation: $\Gamma(E) \subseteq \Gamma(E \cup \hat{F}), \forall \hat{F} \subseteq \hat{E}$.

Strong Benefit to relaxers: All relaxing agents are guaranteed to participate in any maximum matching: $X(\hat{F}) \subseteq \Gamma(E \cup \hat{F}), \forall \hat{F} \subseteq \hat{E}$.

Weak No Harm: $\Gamma(E) \subseteq \Gamma(E \cup \hat{E})$

Weak Benefit to relaxers: $X(\hat{E}) \subseteq \Gamma(E \cup \hat{E})$

Based on the above, we define the following three guarantee combinations:

Strong No Harm, and Strong Benefit to relaxers (SNH-SB): Participation is guaranteed even if some agents do not follow the facilitator's advice. Using this guarantee is preferred when it is expected that several agents will not comply.

Weak No Harm, and Weak Benefit to relaxers (WNH-WB): participation is guaranteed, but only under the assumption that all agents follow the facilitator's advice. This guarantee is preferred when all agents are known to comply as it often results in a larger allocation.

Strong No Harm, and Weak Benefit to relaxers (SNH-WB): the No-harm guarantee holds even with partial compliance, but benefit to relaxers is only guaranteed if all agents comply. In the intermediate case, where we expect very high compliance, the facilitator should weigh the costs of losing the Strong No Harm and/or the strong Benefit to relaxers guarantees and the dissatisfaction of some agents against the benefits of increasing the allocation size and choose one of the three guarantees (See Section 6 for details).

It is easy to see that SNH-SB $\Rightarrow$ SNH-WB $\Rightarrow$ WNH-WB. The example in Figure 2 shows that the hierarchy is strict.

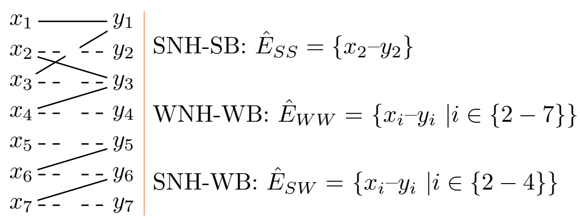

Figure 2. $\hat{E}_{SS}$, $\hat{E}_{WW}$ and $\hat{E}_{SW}$ are relaxation sets for SNH-SB, WNH-WB and SNH-WB. Strictness follows from the fact that $\hat{E}_{SS} \subset \hat{E}_{SW} \subset \hat{E}_{WW}$.

3.3 Problem Statement

Given the above definitions, the facilitator's optimization problem is the following:

Definition 1. *Given X, Y, E, ρ, a participation guarantee PG $\in$ $\{$WNH-WB, SNH-WB, SNH-SB$\}$, aggregated cost bound β, and aggregation function $g \in \{T(\cdot), S(\cdot)\}$, find a relaxation $\hat{E}$ that is PG, $g(\hat{E}) \le \beta$, and $\mu(E \cup \hat{E})$ is the maximum among all such relaxations.*

All in all, with three possible participation guarantees, and two aggregation functions, the above defines six optimization problems. We use the term "solution" to such an optimization problem to refer to a set of relaxations.

3.4 An example

In Figure 1 we match eight courses with six available classrooms. Initially, the course x_2 is compatible with the classroom y_1, x_3 with y_2, x_6 with y_4 and x_7 with y_5. The discomfort levels can vary widely; low for minor inconveniences such as the classroom being far from the offices of the tutor (x_1 and y_1) to major restrictions such as inadequate capacity, in which case, the tutor might have to make necessary arrangements to accommodate extra students (x_4 and y_3). There can be incompatibilities with infinite discomfort. For example, students might require hearing accessories, which might be absent in the classroom (x_3 and y_1). Figure 1 demonstrates the difference between the bound types and the different guarantees. Figure 2 highlights the differences between various participation guarantees.

4 Algorithmic Solutions

Here we introduce efficient algorithms for the different guarantees. In Section 4.1, we present an algorithm for the strong no harm strong benefit guarantee (SNH-SB). In Section 4.2, we demonstrate how to adapt this algorithm for the strong no harm weak benefit guarantee (SNH-WB). For space reasons, some proofs for the SNH-WB algorithm, as well as the algorithmic solution for WNH-WB, are deferred to the full version of the paper [32].

4.1 Strong no harm strong benefit to relaxers (SNH-SB)

We first consider the stronger (and more complex) Strong no harm strong benefit to relaxers guarantee (SNH-SB). We provide an efficient algorithm (Algorithm 1) for the problem and prove its correctness. The core of the algorithm is the weight assignment and the maximum weighted matching calculation (lines 6-8). By assigning appropriate weights to the edges, we ensure that the computed relaxation will be SNH-SB and will have the maximum number of allocations among SNH-SB pairs of the graph G_k. On the loop in lines 3-13, the algorithm adds k dummy agents and connects them to all resources of Y. The dummy agents take the places of other agents in the weighted matching and this way bound the total cost for the non-dummy edges. The algorithm picks the smallest k for which the total cost of a maximum weighted matching (without the dummy edges) does not exceed the bound and removes the added agents for constructing a solution. We will show that the algorithm returns an optimal solution.

It is important to note that the weights w defined in Algorithm 1 are only used by the facilitator to determine the relaxation $\hat{E}$. The allocation itself is performed by a third party *with no weights*. Thus, the facilitator needs to devise appropriate weights that will guarantee the desired properties in the subsequent *weightless* allocation. Devising these weight functions and proving their validity for various participation guarantees is the topic of subsequent sections.

The functions $u(\cdot)$. We define the function $u(\cdot)$ differently for each aggregation function. For the size function, $S(\cdot)$, $u_1(e) \equiv 1$ has value 1 for each relaxation e. In this way, each relaxed edge has a smaller weight compared to the weight of an original edge in the maximum weighted matching computed in line 8 of Algorithm 1. For the total cost function $T(\cdot)$, we define $u_2(e)$ to be the discomfort level of relaxing e. Summing up the contributions of all the $u_2(e)$ values to the weight function, results in a factor that is equal to the total cost, divided by some constant.

Algorithm 1 SNH-SB Optimization

Input: Given X, Y, E, E_R, ρ, aggregated cost bound β, aggregation function $g \in \{T(\cdot), S(\cdot)\}$ and weight function $w : E \cup E_R \to R^+$

Output: A relaxable set $\hat{E}$

1: $G \leftarrow (X, Y, E \cup E_R)$
2: $k_{low} \leftarrow 0, k_{high} \leftarrow |Y|$
3: **while** $k_{low} \leq k_{high}$ **do**
4: $\quad k = \left\lfloor \frac{k_{\text{low}} + k_{\text{high}}}{2} \right\rfloor$
5: $\quad$ Construct G_k by adding k dummy agents to G and connecting them to all resources of Y with edges E_k
6: $\quad$ **for all** $e \in E_k$ **do**
7: $\quad\quad w(e) = 1 + \sum_{e \in E \cup E_R} w(e)$.
8: $\quad$ Compute a maximum weighted matching M_k of G_k
9: $\quad$ **if** $g(M_k \setminus (E \cup E_k)) \leq \beta$ **then**
10: $\quad\quad E_{min} \leftarrow E_k$, $M_{min} \leftarrow M_k$
11: $\quad\quad k_{high} \leftarrow k - 1$
12: $\quad$ **else**
13: $\quad\quad k_{low} \leftarrow k + 1$
14: $\hat{E} = M_{min} \setminus (E \cup E_{min})$
15: **return** $\hat{E}$

We define the weight function w as:

$$w(e) = \begin{cases} (|X| + 1)^2 & e \in E \\ |X| + 1 - \frac{u(e)}{\max_{e' \in E_R}\{u(e')\}} & e \in E_R \end{cases}$$

The main motivation behind this definition is to assign high weights to the original edges, ensuring that the maximum matching of the graph includes them as a subgraph even after any relaxations. Building on this, the function aims to increase the overall size of the matching by incorporating relaxable edges according to their costs.

We have the following result for Algorithm 1.

Theorem 1. *Algorithm 1 with the weights function $w(\cdot)$ solves the SNH-SB optimization problem with aggregate function $g \in \{T(\cdot), S(\cdot)\}$ and bound β.*

The runtime of one iteration of the loop, starting in line 3 of Algorithm 1, is dominated by the weighted matching calculation, which can be done in $O(\max(|X|, |Y|)^3)$ time using the Hungarian algorithm. The loop repeats $O(\log(|Y|)$ times and therefore the total running time is $O(\log(|Y|) * \max(|X|, |Y|)^3)$.

We now outline the proof of the theorem. The full proof appears in an extended version of this paper [32]. Recall that for an edge set F, $\mu(F)$ is the size of a maximum allocation of F. Recall also the definition of a minimal relaxation (Section 3.1). To provide the strong no harm strong benefit guarantee (SNH-SB), we must provide two participation guarantees, namely no-harm and benefit to relaxers, even when some agents do not follow the facilitator's advice. From definitions of the SNH-SB guarantees, it follows that for each SNH-SB relaxation $\hat{E}$, any subset of $\hat{E}$ also provides the SNH-SB guarantees. Therefore, any minimal relaxation $\hat{E}' \subseteq \hat{E}$ for which $\mu(\hat{E}') = \mu(\hat{E})$ should also be SNH-SB. Therefore, it suffices for Algorithm 1 to consider relaxations which are minimal. We will show in Lemma 2 that for all minimal SNH-SB relaxations, $\mu(E \cup \hat{E}) = \mu(E) + |\hat{E}|$. Therefore, any maximum matching of $(X, Y, E \cup \hat{E})$ must contain a maximum matching of (X, Y, E) as a subgraph. We filter out all solutions without these properties by modifying the weight function and assigning significantly greater weights to the E edges.

We show in Lemma 1 that Algorithm 1 indeed filters out all non SNH-SB relaxations (see the proofs of Lemma 1(1) and (2)) and that the returned relaxation maximizes the allocation size (see Lemmas 1(3) and (4) and also Lemma 7 in the extended version of the paper [32]). In the following lemmas, we denote $k_{\min}$ as the value of k, used in line 14 to define $M_{\min}$ and E_{min}. The values of $M_{\min}$ and E_{min} in line 14 are denoted as $M_{k_{\min}}$ and $E_{k_{min}}$. We define a precedence order $\preceq$ on subsets of E_R. Let $\hat{E}, \hat{F} \subseteq E_R$ be two relaxation sets. We say that $\hat{E} \preceq \hat{F}$ if and only if $\hat{E}$ appears before $\hat{F}$ in the order. Let $u : E_R \to \mathbb{R}$. We say that u *represents* $\preceq$ if $\hat{F} \preceq \hat{E} \iff \sum_{e \in \hat{F}} u(e) \leq \sum_{e \in \hat{E}} u(e)$. To prove lemma 1, we will use Lemmas 2 and 3 which are defined later.

Lemma 1. *The set, $\hat{E}$ returned by Algorithm 1 has the following properties:*

(1) It provides the strong benefit for relaxers.
(2) It provides the strong no harm guarantee.
(3) For any SNH-SB set $\hat{F} \subseteq E_R$, $|M_{k_{min}}| \geq \mu(E \cup E_{k_{min}} \cup \hat{F})$.
(4) For any SNH-SB set $\hat{F} \subseteq E_R$ and for any aggregation function g, if $|M_{k_{min}}| = \mu(E \cup E_{k_{min}} \cup \hat{F})$, then $\hat{E} \preceq \hat{F}$.

Set $w^* = |X| + 1$ and let $\hat{E} = M_{k_{min}} \setminus E$ be the relaxation returned by the algorithm.

Proof idea of lemmas 1(1) and 1(2). Consider $\hat{F} \subseteq \hat{E}$. Let $\bar{M} = M_{k_{min}} \cap E$. We will first show that $|\bar{M}_{\hat{F}}| = k_{min} + \mu(E) + |\hat{F}|$. So, (i) $|\bar{M}_{\hat{F}} \cap E| = \mu(E)$, and (ii) $|\bar{M}_{\hat{F}} \cap \hat{F}| = |\hat{F}|$. So, by (i) $\bar{M}_{\hat{F}} \cap E$ is a maximum matching of (X, Y, E), so $\Gamma(E) \subseteq \bar{M}_{\hat{F}}$, and by (ii) $\hat{F} \subseteq \bar{M}_{\hat{F}}$. Since it is correct for any $\hat{F}$, $\hat{E}$ has the no-harm and benefit for relaxers SNH-SB guarantees with respect to $(X \cup X_{k_{min}}, Y, E \cup E_{k_{min}} \cup \hat{E})$.

By construction of $E_{k_{min}}$, all agents $X_{k_{min}}$ participate in all maximum matchings of $(X \cup X_{k_{min}}, Y, E \cup E_{k_{min}} \cup \hat{F})$. It can be shown that $\hat{E}$ has also the SNH-SB guarantees also with respect to $(X, Y, E \cup \hat{E})$ as requested. □

Proof idea of lemma 1(3). For purposes of contradiction, suppose that there is another SNH-SB set $\hat{F}$, for which the maximum matching of $(X \cup X_{k_{min}}, Y, E \cup \hat{F} \cup E_{k_{min}})$ is of greater size.

We set $\hat{F}' \subseteq \hat{F}$ as an inclusion minimal subset of $\hat{F}$ for which $\mu(\hat{F}') = \mu(\hat{F})$. In particular, since $\hat{F}'$ is inclusion minimal, $\mu(E \cup E_{k_{min}} \cup \hat{F}) > \mu(E \cup E_{k_{min}} \cup \hat{F} \setminus \{e\}))$. So, the conditions of Lemma 2 hold for $(E \cup E_{k_{min}})$ and $\hat{F}'$. Let $\bar{M}_{\hat{F}'}$ be a maximum matching of $(X, Y, E \cup E_{k_{min}} \cup \hat{F})$. So, since $M_{k_{min}}$ is a matching of $(X \cup X_{k_{min}}, Y, E \cup E_{k_{min}} \cup \hat{E})$, $|\bar{M}_{\hat{F}'}| > |M_{k_{min}}|$. So,

$$\begin{aligned}
|\bar{M}_{\hat{F}'} \cap \hat{F}'| &= |\bar{M}_{\hat{F}'} \setminus E| = |\bar{M}_{\hat{F}'}| - |\bar{M}_{\hat{F}'} \cap E| \\
&= |\bar{M}_{\hat{F}'}| - \mu(E) \;\; \text{by Lemma 2} \\
&> |M_{k_{min}}| - \mu(E) \;\; \text{by assumption} \\
&= |M_{k_{min}}| - |M_{k_{min}} \cap E| \;\; \text{by Lemma 3} \\
&= |M_{k_{min}} \cap \hat{E}| + |M_{k_{min}} \cap E_{k_{min}}|
\end{aligned}$$

Using this equation, it can be shown that $w(M_{k_{min}}) < w(\bar{M}_{\hat{F}'})$. This contradicts the maximality of $w(M_{k_{min}})$. □

Proof of lemma 1(4). Let $\hat{F}$ be a SNH-SB relaxation with maximum allocation when adding k_{min} dummy agents and for which $g(\hat{F}) \leq \beta$.

First note that since $\hat{F}$ is SNH-SB so is any subset of $\hat{F}$. So, for any $e \in \hat{F}$, $\mu(E \cup \hat{F}) > \mu(E \cup \hat{F} \setminus \{e\})$, or else $\hat{F} \setminus \{e\}$ is a SNH-SB set with allocation of the same size that preceding $\hat{F}$ in $\preceq$ (Since $u(\cdot)$ is strictly positive).

For $e \in E_R$, set $u'(e) = \frac{u(e)}{\max_{e' \in E_R}\{w(e')\}}$. So, u' is also a representation of $\preceq$, and

$$w(e) = \begin{cases} (w^*)^2 & e \in E \\ w^* - u'(e) & e \in E_R \end{cases}$$

We have already established that

$$\mu(E \cup E_{k_{min}} \cup \hat{E}) = \mu(E) + k_{min} + |\hat{E}|.$$

And by Lemma 2

$$\mu(E \cup E_{k_{min}} \cup \hat{F}) = \mu(E) + k_{min} + |\hat{F}|.$$

So, since both $\hat{E}$ and $\hat{F}$ have allocation of the same size, $|\hat{E}| = |\hat{F}|$. By Lemma 3 (for the first equality) and Lemma 2 (for the second) we have (see the extended version of the paper [32] for details):

$$\begin{aligned}
w(M_{k_{min}}) &= \\
&\quad w(M_{k_{min}} \cap E_{k_{min}}) - u'(\hat{E}) + |\hat{E}|w^* + \mu(E) \cdot (w^*)^2 \\
w(\bar{M}_{\hat{F}}) &= w(M_{k_{min}} \cap E_{k_{min}}) - u'(\hat{F}) + |\hat{F}|w^* + \mu(E) \cdot (w^*)^2
\end{aligned}$$

By construction $w(M_{k_{min}}) \geq w(\bar{M}_{\hat{F}})$. So, since $|\hat{E}| = |\hat{F}|$, $u'(\hat{E}) \leq u'(\hat{F})$. So, $\hat{E} \preceq \hat{F}$. □

The two lemmas below are used for proving Lemma 1:

Lemma 2. *Let $\hat{E}$ be a minimal SNH-SB relaxation. Then $\mu(E \cup \hat{E}) = \mu(E) + |\hat{E}|$.*

The following three results are used in proving Lemma 2. In these results, $G = (X, Y, E)$ is a bipartite graph.

(1) Let $x \in \Gamma(E)$. Let (x, y) be an edge in E_R. Then $\mu(E) = \mu(E \cup \{(x, y)\})$.

Proof idea of (1). We first note that since $E \subseteq E \cup \{(x, y)\}$, it follows that $\mu(E) \leq \mu(E \cup \{(x, y)\}$. We assume for contradiction that $\mu(E) < \mu(E \cup \{(x, y)\})$ and show that it contradicts the participation of x in all maximum matchings of $G = (X, Y, E)$. □

(2) Let x be an agent that does not participate in all maximum matchings of G and $(x, y) \in E_R \notin E$. If x participates in all maximum matchings of $(X, Y, E \cup \{(x, y)\})$, then $\mu(E \cup \{(x, y)\}) = \mu(E) + 1$.

Proof of (2). We assume that x participates in all maximum matchings of $(X, Y, E \cup \{(x, y)\})$ and prove that $\mu(E) + 1 = \mu(E \cup \{(x, y)\})$. Since $(X, Y, E \cup \{(x, y)\})$ is supergraph of G, $\mu(E) \leq \mu(E \cup \{(x, y)\})$. On the other hand, from definition of matching, $\mu(E \cup \{(x, y)\}) \leq \mu(E) + 1$. Therefore, either $\mu(E \cup \{(x, y)\}) = \mu(E)$ or $\mu(E \cup \{(x, y)\}) = \mu(E) + 1$. Assume for contradiction that $\mu(E \cup \{(x, y)\}) = \mu(E)$. In this case, all maximum matchings of G are also maximum matchings of $(X, Y, E \cup \{(x, y)\})$. From the lemma's definitions, there is a maximum matching of G in which x does not participate. Since it is also a maximum matching of $(X, Y, E \cup \{(x, y)\})$, x does not participate in all maximum matchings of $(X, Y, E \cup \{(x, y)\})$, a contradiction. □

(3) A node $x \in X$ does not participate in all maximum matchings of G iff there is an even-length alternating path between x and a free node with respect to any maximum matching of G. (This well-known result in matching theory [21]; (A proof of this result and the definitions for alternating paths and free nodes are also given in the extended version of the paper [32].)

We use the above three results to prove Lemma 2.

Proof of Lemma 2. We recall that an edge set $\hat{E}$ is minimal if for any edge $e \in \hat{E}$, $\mu(E \cup \hat{E} \setminus \{e\}) < \mu(E \cup \hat{E})$. The proof is by induction on $|\hat{E}|$. The basis is for $|\hat{E}| = 1$, where the lemma is fulfilled trivially. We now assume that the lemma holds for $|\hat{E}| = q$ and prove it for $|\hat{E}| = q + 1$.

Let $|\hat{E}|$ be a SNH-SB relaxation and assume that for any edge $e \in \hat{E}$, $\mu(E \cup \hat{E} \setminus \{e\}) < \mu(E \cup \hat{E})$. Let x be any agent in $X(\hat{E})$, and let $(x, y) \in \hat{E}$ be the corresponding edge in $\hat{E}$. It follows that:

$$\mu(E \cup \hat{E} \setminus \{(x, y)\}) + 1 = \mu(E \cup \hat{E}) \tag{1}$$

Case 1: The set $\hat{E} \setminus \{(x, y)\}$ is minimal: Here, by induction hypothesis, $\mu(E \cup \hat{E} \setminus \{(x, y)\}) = \mu(E) + q$. Therefore, if we merge this condition with Equation (1), we get $\mu(E \cup \hat{E}) = \mu(E \cup \hat{E} \setminus \{(x, y)\}) + 1 = \mu(E) + q + 1$ as required.

Case 2: The set $\hat{E} \setminus \{(x, y)\}$ is not minimal: In this case, there is an edge $(x', y') \in \hat{E} \setminus \{(x, y)\}$ such that $\mu(E \cup \hat{E} \setminus \{(x, y)\}) = \mu(E \cup \hat{E} \setminus \{(x, y), (x', y')\})$. Let M be a maximum matching of $(X, Y, E \cup \hat{E} \setminus \{(x, y)\})$ without (x', y'). (Such a matching exists because of Equation (1) above.)

We note that since $\mu(E \cup \hat{E}) > \mu(E \cup \hat{E} \setminus \{(x, y)\})$, Result (2) above implies that x must not participate in all maximum matchings of $(X, Y, E \cup \hat{E} \setminus \{(x, y)\})$. Therefore, it follows from Result (3) above that there is an even-length alternating path between x and a free agent concerning any maximum matching of $(X, Y, E \cup \hat{E} \setminus \{(x, y)\})$. On the other hand, since $\hat{E}$ has the SNH-SB benefit to relaxers, x' must participate in all maximum matchings of $(X, Y, E \cup \hat{E} \setminus \{(x, y)\})$. Therefore, by Result (3) above, x' does not appear in any even length alternating path between x and a free vertex in any maximum matching of $(X, Y, E \cup \hat{E} \setminus \{(x, y)\})$.

Since M is also a maximum matching of $(X, Y, E \cup \hat{E} \setminus \{(x, y), (x', y')\})$, an alternating path of even length between x and a free vertex of M exists in $(X, Y, E \cup \hat{E} \setminus \{(x, y), (x', y')\})$ and therefore x does not participate in all maximum matchings of $(X, Y, E \cup \hat{E} \setminus \{(x, y), (x', y')\})$.

We now add (x, y) to $(X, Y, E \cup \hat{E} \setminus \{(x, y), (x', y')\})$ and get $(X, Y, E \cup \hat{E} \setminus \{(x', y')\})$. If x participates in all maximum matchings of $(X, Y, E \cup \hat{E} \setminus \{(x', y')\})$, then by Result (1) above, the matching size must increase, contradicting the fact that the original set is minimal. Otherwise, x does not participate in all maximum matchings after x relaxes its restrictions, contradicting the SNH-SB. Therefore, we conclude that no such x' exists, and we are done. □

Lemma 3 below is based on the definition of $w(\cdot)$.

Lemma 3. $|M_{k_{min}} \cap E| = \mu(E)$.

4.2 *Strong no harm weak benefit guarantee (SNH-WB)*

As in the SNH-SB problem, the SNH-WB problem can also be solved using Algorithm 1. Intuitively, as in the right column of Figure 1, it is possible that a maximum matching of $(X, Y, E \cup \hat{E})$ does not have a maximum matching of (X, Y, E) as a subgraph. Therefore, the weights of the E edges must be reduced. However, these weights must remain greater than those of $\hat{E}$ to ensure that the WB guarantee is preserved.

Note that if the set $\hat{E}$ returned by the algorithm contains an edge $\{x, y\} \in \hat{E}$ such that $x \in \Gamma(E)$, then x cannot participate in all maximum matchings of $(X, Y, E \cup \hat{E} \setminus \{x, y\})$ and therefore, the SNH-WB no-harm guarantee does not hold. That is so since otherwise, $\mu(E \cup \hat{E} \setminus \{x, y\}) = \mu(E \cup \hat{E})$ and the maximum weighted matching in line 8 won't contain the relaxed edge $\{x, y\}$, contradicting the fact that $\{x, y\} \in \hat{E}$. Therefore, we assign a negative weight to all edges with agent in $\Gamma(E)$ and this way, we keep the no-harm guarantee even when some agents of $\hat{E}$ do not relax. The updated weight function is defined as follows:

$$w(e) = \begin{cases} -1 & e = (x, y) \in E \cup E_R, x \in \Gamma(E) \\ (|X| + 1) & e = (x, y) \in E, x \notin \Gamma(E) \\ |X| + 1 - \\ \quad \frac{u(e)}{\max_{e' \in E_R}\{u(e')\}} & e = (x, y) \in E_R, x \notin \Gamma(E) \end{cases}$$

In this function, the weights of the original edges are smaller as it is not required to have a maximum matching from original edges as a subgraph of the resulting matching; however, it is required that the guaranteed agents won't be asked to relax their restrictions. This way, if they do not relax, no harm will be caused.

Note that the set $\Gamma(E)$ can be computed in polynomial time using the method of [18]. The functions u that are used are identical to those used in the SNH-SB case. We conclude with the following theorem:

Theorem 2. *Algorithm 1 with the weights function $w(\cdot)$ solves the SNH-WB optimization problem with an aggregate cost function g (see Theorem 1) and a bound β.*

5 Many-to-One and One-to-Many Allocations

Up to now, we have considered allocating a single resource to each agent. In many scenarios like allocation of classrooms to courses [31], a single agent might need multiple resources, e.g., for providing a course with enough classrooms for holding an exam where students need to be far from each other. In other scenarios, a single resource might be shared among multiple agents, e.g., for providing multiple courses with very few participants with an appropriate classroom for holding an exam together. This section formulates an optimization problem for many-to-one allocations and presents a polynomial time algorithm to solve it. The one-to-many case where a resource can be shared by multiple agents can be done similarly.

The problem is similar to Definition 1 but with an addition: a function $d : X \to \mathbb{N}$ that assigns to each agent its desired amount of resources is added to the problem inputs. Note that $\mu(\cdot)$ is the size of a maximum valid allocation with respect to d (not necessarily a maximum matching).

Note when moving to the many-to-one setting, the aggregation functions are naturally applied to the set of relaxed edges (not relaxing agents). In particular, the size aggregation cost function bounds the overall number of relaxed edges.

We propose Algorithm 2 that uses Algorithm 1. Algorithm 2 first duplicates the agents that need more than one resource and adds some relevant edges. It then runs Algorithm 1 on the modified graph and returns its results. Lemma 4 is crucial for showing that this modification does not affect the desired participation guarantees.

Lemma 4. *Let $G = (X, Y, E)$ be a graph and let $x, x' \in X$ be agents. Set Y_x and $Y_{x'}$ as the resources adjacent to x and x' in E. If $Y_x \subseteq Y'_x$ and $x \in \Gamma(G)$, then x' is also in $\Gamma(G)$.*

Proof of lemma 4. Assume for the purposes of contradiction that $x' \notin \Gamma(G)$. Let M be a maximum matching of G without x'. Since $x \in \Gamma(G)$, there is y for which $\{x, y\} \in M$. Since $Y_x \subseteq Y'_x$, the edge $\{x', y\} \in E$. therefore, we can substitute the edges and have a maximum matching of G without x, and we have a contradiction. □

Algorithm 2

Input: Given, X, Y, E, ρ, aggregated cost bound β, aggregation function $g \in \{T(\cdot), S(\cdot)\}$, demand function $d : X \to N$ and a weight function $w : E \cup E_R \to R^+$

Output: A relaxable set $\hat{E}$

1: $X', Y', E' \leftarrow \emptyset$
2: **for all** agents $x \in X$ **do**
3: $\quad X' = X' \cup \{x_j | 0 \leq j < d(x)\}$.
4: **for all** resources $y \in Y$ **do**
5: $\quad Y' = Y' \cup \{y\}$.
6: $E' = \{\{x_j, y\} | \{x, y\} \in E\}$,
7: $E'_R = \{\{x_j, y\} | \{x, y\} \in E_R\}$,
8: $\forall \{x_j, y\} \in E'_R, \rho(\{\{x_j, y\}\}) = \rho(\{x, y\})$,
9: $\forall \{x_j, y\} \in E' \cup E'_R, w'(x_j, y) = w(x, y)$
10: Run Algorithm 1 with $X', Y', E', E'_R, \rho', \beta, g$ and w'. Let $\hat{E}'$ be the returned edge set.
11: Construct a set $\hat{E}$ from $\hat{E}'$ by substituting all instances x_j of x by x itself.
12: **return** $\hat{E}$

We conclude with the following:

Theorem 3. *Algorithm 2 with a weight function $w(\cdot)$ corresponds to a participation guarantee PG $\in$ {SNH-SB, WNH-WB, SNH-WB} (e.g., the weight functions in Section 4) solves the many-to-one PG optimization problem with an aggregate cost function g (see Theorem 1), a bound β and a demand function $d(\cdot)$.*

Proof of Theorem 3. Note that Algorithm 2 uses Algorithm 1 as a subprocedure. The correctness of Algorithm 1 was proved in the previous sections. We conclude that Algorithm 1 returns a relaxation set with a maximum sized allocation that has the relevant participation guarantee with the relevant aggregate function and bound. It remains to show that relevant participation guarantee is preserved when merging the duplicated agents as we do in line 11 of Algorithm 2.

Lemma 4 implies that if one instance of an agent is guaranteed to be matched so is for all other instances. Therefore, the different guarantees for relaxers and the different no-harm guarantees provided by Algorithm 1 to some instances of an agent hold for all and therefore for the agent itself. This concludes the proof of Theorem 3. □

The runtime of Algorithm 2 is determined by the runtime of Algorithm 1 on line 10. The key difference lies in the size of X', which is $\sum_{x \in X} d(x)$. The overall complexity is therefore $O(\log(|Y|) \cdot \max(\sum_{x \in X} d(x), |Y|)^3)$.

6 Experiments

In this section, we discuss the results of employing a facilitator, as described in this paper, on three different real-world datasets. Experiments were conducted in all three allocations settings (one-to-one, one-to-many and many-to-one). We compared the WNH-WB, SNH-WB and the SNH-SB participation guarantees and the extent to which these different participation guarantees affect the resultant allocation size.[2].

[2] The code and data for running the experiments are available at [29]

Datasets. We ran our experiments on three real-world datasets. The **courses and classes dataset (COURSE)** [30] contains 154 courses and 144 classrooms, with the goal of matching courses to appropriate classrooms. The attributes in this dataset are room capacity, location, availability of accommodations for physical disability, and availability of accessories for hearing disability. The **students lab dataset (LAB)** [31] consists of a lab with 31 students that should be matched to seats in 14 lab rooms. Room attributes include proximity to some locations in the lab (e.g., advisor's room, restrooms, kitchen), capacity, and strength of the Wi-Fi signal. The **children activities dataset (CHILD)** [33] consists of 653 children and 533 available activities over a period of several weeks. We focused on one week during the autumn vacation in the Swiss municipality of Morges. During that week, 249 activities were offered. Attributes included minimum and maximum allowed age and children's priorities over the activities. Note that the CHILD dataset has significantly more agents and relatively less available resources, so the overall match rate and the potential for improvements by the facilitator are much more limited.

Appropriate **discomfort** functions were defined for the datasets. In the LAB dataset, the attributes were explicitly rated on a scale of 1–5, and the discomfort function was defined accordingly, with rating 1 taken as no discomfort, and 2-5 having discomfort of 1-4, respectively. The discomfort for relaxing an edge is the sum of the discomforts for all its attributes. For the COURSE dataset, the discomfort was estimated from the nature of the constraints. These estimates were informed by consultations with a university administrator and faculty members. The availability of accommodations for physical disability and the availability of accessories for hearing disability have a crucial impact on the discomfort level. The distance from the desired location determines the discomfort incurred by the room's location. Finally, the discomfort of the room capacity is more significant as the difference from the desired capacity is more significant.

The CHILD dataset has attributes of both types (explicit rating and not), so both methods were employed for determining the discomfort function. Children rated some activities, and their ratings contributed to the discomfort in a way similar to the attributes in the LAB dataset. If a child is younger than the minimum age for the activity, the difference between her age and the minimum age of the activity was considered for determining the discomfort. The contributions to the discomfort when the child is older than the maximum allowed age for the activity were computed similarly.

Maximum possible match sizes when all agents comply. Here, we compare the WNH-WB, SNH-WB and SNH-SB relaxations for the special case where all agents comply. Since a SNH-SB relaxation is optimized to handle the case when not all agents comply, this stricter condition might impact the matching size when compared with WNH-WB, even when all agents comply. The SNH-WB is somewhere between the SNH-SB and the WNH-WB in the courses dataset and is comparable with SNH-SB in the two other datasets. Figure 3 depicts the resultant matching sizes for WNH-WB, SNH-WB and SNH-SB relaxations when all agents comply. The baseline bar is for when no relaxation is allowed. First, note that relaxation does indeed increase the match size considerably, even for low bounds. In the CHILD dataset, a 5% increase in the relaxation cost bound allows 32 more children to be matched. Comparing the results of WNH-WB, SNH-WB and SNH-SB, we see that while SNH-SB and SNH-WB do provide better results sometimes in smaller matchings, the difference is relatively small.

Match sizes when not all agents comply. Figure 4 depicts the matching size as a function of the number of complying agents, with

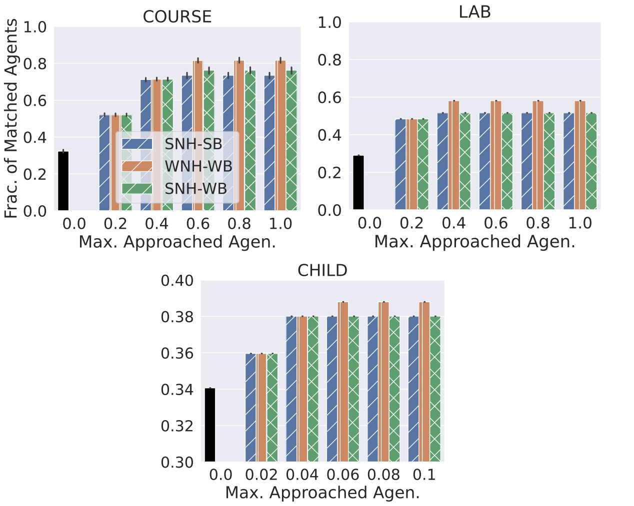

Figure 3. Fraction of matched agents as a function of the bound on allowed relaxations, for the SNH-SB, SNH-WB and WNH-WB guarantees. The scale and limits of the y axis are different for the different datasets. The baseline (in black) corresponds to no relaxations or a cost bound of zero.

no bound restrictions. Here, we first computed the full SNH-SB, SNH-WB and WNH-WB relaxations, using the methods devised in the earlier sections, randomly selected a set of complying agents (of different sizes), and then computed the maximum matching with the resultant relaxation. Each point on the graph is an average of ten such random samples. We can see that SNH-SB performs better with low and moderate compliance rates, while WNH-WB performs better when more agents comply. The performance of SNH-WB is comparable to that of SNH-SB in most cases and it is between SNH-SB and WNH-WB for the COURSE dataset.

Hence, we propose using SNH-SB when it is expected that several agents will not comply. However, if all agents are known to comply, WNH-WB is preferable. In the intermediate case, where we expect very high compliance, the facilitator should weigh the costs of losing the SNH-SB or the SNH-WB guarantees and the dissatisfaction of some agents against the benefits of increasing the allocation size.

Figure 4. Fraction of matched agents vs. number of relaxing agents for SNH-SB, SNH-WB, and WNH-WB guarantees. The scale and limits of the y axis are different for the different datasets.

Many-to-one and one-to-many. For the COURSE dataset, we allow some courses to have multiple rooms. Such a scenario may happen in exams where there needs to be some gap between students for preventing dishonesty. The CHILD dataset is used when some of the activities allow multiple children for fixing the balance between the children and the avaialabe activities. For the CHILD dataset, we chose 1-3 instances at random for each activity. For the COURSE dataset, we again chose 1-3 instances at random for each course. Similar to Figure 4, In Figures 5 (left) we can see that SNH-SB performs better with low and moderate compliance rates, while WNH-WB performs better when more agents comply. The bumps in the graphs, where the number of relaxing agents is greatest, is due to the fact that the facilitator only addresses so many agents in a few of the replications and we can therefore safely ignore these bumps. In Figure 5 (right) (COURSE), the fraction of matches is lower than in Figure 3. That is since some courses might need more than one classroom so there is stronger competition on the available classrooms. The opposite is true for the CHILD dataset, where multiple children are allowed per activity.

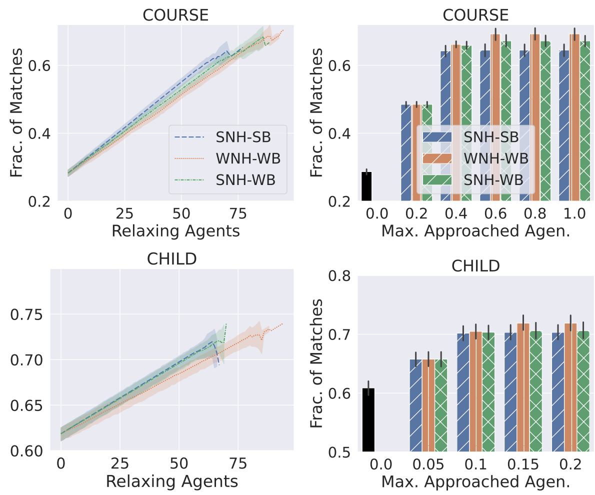

Figure 5. Results for many to one (COURSE) and one to many (CHILD)

7 Conclusion

In this work, we define a matching problem where a facilitator advises agents on which constraints to relax. The facilitator's goal is to increase the allocation size while ensuring that no agent is harmed by the suggested relaxations and that relaxing agents benefit from relaxing by securing a guaranteed resource. Additionally, the facilitator is constrained by a bound on the total cost of relaxations that can be suggested.

For three allocation settings, namely one-to-one, many-to-one and one-to-many, we consider three variants of participation guarantees and two possible forms of cost aggregration. A general approach is presented to tackle the one-to-one problems leading to polynomial time algorithms. An extension of this approach is presented to obtain polynomial time algorithms for the many-to-one and one-to-many problems as well. Our experiments demonstrate the usefulness of these algorithms by applying them to three real-world datasets.

Expanding our research by developing other techniques for defining discomfort levels and quantifying them through numerical values is left for future work.

Acknowledgment: We thank the ECAI-2025 reviewers for carefully reading the paper and providing very helpful comments. *The research of Yonatan Aumann is supported in part by ISF grant 3007/24. The research of Sarit Kraus is supported in part by ISF grant 2544/24.

References

[1] HPD Housing Ambassadors. https://www.nyc.gov/site/hpd/services-and-information/housing-ambassadors.page. Accessed: 2025-05-04.

[2] Housing Connect 2 | Public Portal. https://housingconnect.nyc.gov. Accessed: 2025-05-04.

[3] J. C. Aguma. A matching mechanism for provision of housing to the marginalized. In *Proc. SAI*, pages 20–33. Springer, 2022.

[4] Y. Bao, G. Zang, H. Yang, Z. Gao, and J. Long. Mathematical modeling of the platform assignment problem in a ride-sourcing market with a third-party integrator. *Transportation Research Part B: Methodological*, 178:102833, 2023.

[5] F. Bobbio, M. Carvalho, A. Lodi, I. Rios, and A. Torrico. Capacity planning in stable matching: An application to school choice. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pages 295–295, 2023.

[6] N. Boehmer, R. Bredereck, K. Heeger, and R. Niedermeier. Bribery and control in stable marriage. *Journal of Artificial Intelligence Research (JAIR)*, 71:993–1048, 2021.

[7] H. Chan, L. Tran-Thanh, B. Wilder, E. Rice, P. Vayanos, and M. Tambe. Utilizing housing resources for homeless youth through the lens of multiple multi-dimensional knapsacks. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 41–47, 2018.

[8] J. Chen and G. Csáji. Optimal capacity modification for many-to-one matching problems. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, pages 2880–2882, 2023.

[9] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaitre, N. Maudet, J. Padget, S. Phelps, J. A. Rodrígues-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica (Slovenia)*, 30(1): 3–31, 2006.

[10] M.-C. Costa. Persistency in maximum cardinality bipartite matchings. *Operations Research Letters (ORL)*, 15(3):143–149, 1994.

[11] D. A. Dolgov and E. H. Durfee. Resource allocation among agents with MDP-induced preferences. *Journal of Artificial Intelligence Research (JAIR)*, 27:505–549, 2006.

[12] A. L. Dulmage and N. S. Mendelsohn. Coverings of bipartite graphs. *Canadian Journal of Mathematics*, 10:517–534, 1958.

[13] P. Faliszewski and J. Rothe. Control and bribery in voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, page 146–168. Cambridge University Press, New York, NY, 2016.

[14] F. O. Frimpong and A. Owusu. Allocation of classroom space using linear programming (A case study: Premier Nurses Training College, Kumasi). *Journal of Economics and sustainable Development*, 6(2): 12–20, 2015.

[15] S. Gokhale, S. Singla, S. Narang, and R. Vaish. Capacity modification in the stable matching problem. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, pages 697–705, 2024.

[16] V. Gorodetski, O. Karsaev, and V. Konushy. Multi-agent system for resource allocation and scheduling. In *Proc. CEEMAS*, pages 236–246, Heidelberg, Germany, 2003. Springer.

[17] G. Huang, D. Li, S. T. Ng, L. Wang, and T. Wang. A methodology for assessing supply-demand matching of smart government services from citizens' perspective: A case study in Nanjing, China. *Habitat International*, 138:102880, 2023.

[18] R. W. Irving, T. Kavitha, K. Mehlhorn, D. Michail, and K. E. Paluch. Rank-maximal matchings. *ACM Transactions on Algorithms (TALG)*, 2 (4):602–610, 2006.

[19] V. Kilambi, M. Barah, R. N. Formica, J. J. Friedewald, and S. Mehrotra. Evaluation of opening offers early for deceased donor kidneys at risk of nonutilization. *Clinical Journal of the American Society of Nephrology*, pages 10–2215, 2023.

[20] M. Lin, L. Ma, and C. Ying. Matching daily home health-care demands with supply in service-sharing platforms. *Transportation Research Part E: Logistics and Transportation Review*, 145:102177, 2021.

[21] L. Lovász and M. D. Plummer. *Matching theory*. North Holland, Amsterdam, The Netherlands, 2016.

[22] T. T. Nguyen, M. Roos, and J. Rothe. A survey of approximability and inapproximability results for social welfare optimization in multiagent resource allocation. *Annals of Mathematics and Artificial Intelligence*, 68(1-3):65–90, 2013.

[23] P. Pan, A. Jin, A. Mahmoudi, X. Li, and Y. Wang. A matching model for construction subcontractor selection in engineering bid decisions using ordinal priority approach. *Journal of Asian Architecture and Building Engineering*, pages 1–14, 2023.

[24] P. S. Rao, D. E. Schaubel, M. K. Guidinger, K. A. Andreoni, R. A. Wolfe, R. M. Merion, F. K. Port, and R. S. Sung. A comprehensive risk quantification score for deceased donor kidneys: The kidney donor risk index. *Transplantation*, 88(2):231–236, 2009.

[25] A. Sharam, M. C. Byford, B. Karabay, S. McNelis, and T. Burke. Matching markets in housing and housing assistance. AHURI Final Report No. 307, 2018.

[26] V. W. Slaugh, M. Akan, O. Kesten, and M. U. Ünver. The Pennsylvania adoption exchange improves its matching process. *Interfaces*, 46(2): 133–153, 2016. URL http://www.jstor.org/stable/45154095.

[27] G. Slingerland, I. Mulder, and T. Jaskiewicz. Empowering community volunteers through matchmaking services. In *Linköping Electronic Conference Proceedings*, 2018.

[28] Z. Tang, Y. Jiang, and F. Yang. An efficient Lagrangian relaxation algorithm for the shared parking problem. *Computers & Industrial Engineering*, 176:108860, 2023.

[29] Y. Trabelsi. Code and data: Facilitating matches on allocation platforms. https://github.com/yohayt/Facilitating-Matches-on-Allocation-Platforms, 2025. Accessed: 24/08/2025.

[30] Y. Trabelsi, A. Adiga, S. Kraus, and S. S. Ravi. Resource allocation to agents with restrictions: Maximizing likelihood with minimum compromise. In *European Conference on Multi-Agent Systems*, pages 403–420, Heidelberg, Germany, 2022. Springer.

[31] Y. Trabelsi, A. Adiga, S. Kraus, S. S. Ravi, and D. J. Rosenkrantz. Resource sharing through multi-round matchings. In *Proc. 37th AAAI*, pages 11681–11690. AAAI Press, 2023.

[32] Y. Trabelsi, A. Adiga, Y. Aumann, S. Kraus, and S. S. Ravi. Facilitating matches on allocation platforms. *arXiv preprint arXiv:2508.18325*, 2025.

[33] S. Varone and C. Beffa. Dataset on a problem of assigning activities to children, with various optimization constraints. *Data in Brief*, 25: 104168, 2019.

[34] J. Yu, Y. Fang, Y. Zhong, X. Zhang, and R. Zhang. Pricing and quality strategies for an on-demand housekeeping platform with customer-intensive services. *Transportation Research Part E: Logistics and Transportation Review*, 164:102760, 2022.

[35] Z. Zahedi, S. Sengupta, and S. Kambhampati. 'Why didn't you allocate this task to them?' Negotiation-aware explicable task allocation and contrastive explanation generation. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, pages 2292–2294. IFAAMAS, 2023.

[36] Z. Zahedi, S. Sengupta, and S. Kambhampati. 'Why didn't you allocate this task to them?' Negotiation-aware task allocation and contrastive explanation generation. In *Proc. 34th AAAI*, pages 10243–10251, 2024.

[37] X. Zhang, J. Han, and W. Zhang. An efficient algorithm for finding all possible input nodes for controlling complex networks. *Scientific Reports*, 7(1):10677, 2017.