Achieving Balanced Representation in School Choice with Diversity Goals

Zhaohong Sun[1, 2], **Makoto Yokoo**[1]

[1]Kyushu University, Japan
[2]AI Lab, CyberAgent, Japan

Abstract

Student placements under diversity constraints are a common practice globally. This paper addresses the selection of students by a single school under a one-to-one convention, where students can belong to multiple types but are counted only once based on one type. While existing algorithms in economics and computer science aim to help schools meet diversity goals and priorities, we demonstrate that these methods can result in significant imbalances among students with different type combinations.

To address this issue, we introduce a new property called balanced representation, which ensures fair representation across all types and type combinations. We propose a straightforward choice function that uniquely satisfies four fundamental properties: maximal diversity, non-wastefulness, justified envy-freeness, and balanced representation. While previous research has primarily focused on algorithms based on bipartite graphs, we take a different approach by utilizing flow networks. This method provides a more compact formalization of the problem and significantly improves computational efficiency. Additionally, we present efficient algorithms for implementing our choice function within both the bipartite graph and flow network frameworks.

Introduction

School choice programs with diversity considerations are widespread globally and have been extensively studied in the economics and computer science literature (Biró et al. 2010; Echenique and Yenmez 2015; Kurata et al. 2015; Goto et al. 2016; Kamada and Kojima 2018; Aziz et al. 2019; Baswana et al. 2019). In these programs, students are often associated with multiple types, such as gender, race, disability, or special talent, and schools typically establish quotas or targets for each type to achieve diverse and balanced outcomes.

Many studies implicitly adopt a *one-to-all convention*, where each student is counted under every type they belong to, as observed in Israeli gap year programs (Gonczarowski et al. 2019). In contrast, the *one-to-one convention* is more practical in settings like college admissions in India (Sönmez and Yenmez 2022) and school choice in Chile (Correa et al. 2022). Under this convention, each student is counted under only one of their associated types. For example, an Aboriginal girl might occupy either a seat reserved for Aboriginal students or one reserved for girls, but not both. The one-to-one convention also applies in other critical contexts, such as vaccine allocation (Pathak et al. 2021; Aziz and Brandl 2021).

The one-to-one convention has been extensively studied in the school choice literature, with most studies assuming that students have strict preferences over reserved seats of different types (or resolving ties to create a strict preference ordering) (Aygün and Turhan 2017; Kurata et al. 2017; Baswana et al. 2019). However, students are often indifferent to the type of reserved seat they are assigned, provided they secure a spot at their desired school. Imposing strict preferences over reserved types can lead to unintended consequences, potentially impacting the assignments of other students. Additionally, this approach does not always optimally fulfill diversity goals. To address these limitations, Sönmez and Yenmez (2022) proposed a novel method called "smart reserves" where all reserved seats aimed at achieving diversity goals are maximally satisfied without requiring students to express preferences over different types of reserved seats. Aziz and Sun (2021a, 2025) further extended this concept to a more general framework with "multiple ranks of reserves," accommodating a broader spectrum of diversity constraints such as minimum and maximum quotas and proportionality requirements.

An important issue that has not been adequately addressed in previous work is the potential for imbalanced outcomes among students with distinct type combinations. For instance, consider a school that reserves seats for both aborigines and girls. Existing algorithms might lead to scenarios where the seats reserved for aborigines are filled exclusively by aboriginal boys, while the seats reserved for girls are occupied by non-aboriginal students. Consequently, no aboriginal girls are matched, even though the quotas for both types are technically met. Such an outcome fails to achieve the original objective of fostering a diverse and balanced integration of students from different backgrounds. Ideally, the goal is to compute an assignment that ensures fair representation not only for each individual type but also for combinations of types.

Another critical issue is computational efficiency. This concern becomes particularly relevant because, in practice, the number of privilege types is typically small, leading to a significantly smaller number of distinct type combinations compared to the number of students. For instance, in Brazilian college admissions, there are three privilege types: public high school, low-income, and racial minority (Aygün and Bó 2020). Similarly, in Indian college admissions, the social privilege groups include Scheduled Castes (SCs), Scheduled Tribes (STs), and Other Backward Classes (OBCs), alongside additional reservation categories such as gender and disability (Sönmez and Yenmez 2022).

Existing selection rules for schools are based on bipartite graph structures that model the relationship between students and reserved seats (Sönmez and Yenmez 2022; Aziz and Sun 2021a). In contrast, we show that the problem can be represented more compactly by leveraging flow networks, where nodes correspond to types and type combinations. This approach offers a key advantage: the associated flow network problem can be solved in strongly polynomial time, making it independent of the number of students.

In this paper, we adopt the model of multiple ranks of quotas and follow the one-to-one matching convention. While maintaining ranked quotas for different types as the primary diversity goals, we aim to address imbalances among students with different type combinations to promote a more equitable and integrated outcome. We summarize our contributions as follows:

- First, we introduce a new property called *balanced representation* to address the issue of imbalance. This property aims to maximize the minimum selection ratio across all matchings that adhere to maximal diversity.
- Second, we propose a straightforward choice function for schools that *uniquely* satisfies four key properties: i) maximal diversity, ii) non-wastefulness, iii) justified envy-freeness, and iv) balanced representation.
- Third, we develop a novel graph structure based on flow networks, providing a more compact representation of the problem compared to existing methods.
- Fourth, we present efficient algorithms to implement the newly proposed choice function using two graph structures: ranked reservation graphs, as studied in previous work, and flow networks, which we propose.

Related Work

Abdulkadiroğlu and Sönmez (2003) and Abdulkadiroğlu (2005) explored the school choice model with rigid racial and gender quotas, imposing limits on the number of admitted students from specific groups. Schools commonly set both maximum and minimum quotas for each student type (Hafalir, Yenmez, and Yildirim 2013; Kojima 2012; Kominers and Sönmez 2013). Ehlers et al. (2014) developed the *controlled school choice* model, analyzing the effects of treating diversity quotas as hard or soft bounds, while Echenique and Yenmez (2015) investigated choice functions that satisfy substitutability.

These studies assume that each student belongs to a single type. However, real-world programs in Brazil, Chile, Israel, and India consider overlapping types (Aygün and Turhan 2016; Baswana et al. 2019; Correa et al. 2019; Kurata et al. 2017; Gonczarowski et al. 2019). When individuals have multiple types, two main conventions determine seat allocation: the *one-for-all* and *one-for-one* conventions (Sönmez and Yenmez 2022).

In the one-for-all convention, individuals occupy reserved seats for all qualifying types (Aziz 2019; Aziz, Gaspers, and Sun 2020; Gonczarowski et al. 2019), but optimally meeting diversity goals under this convention is NP-hard (Biró et al. 2010). The one-for-one convention, on the other hand, assigns an individual to a single reserved seat, either based on strict preferences over types (Aygün and Turhan 2016; Kurata et al. 2017), fixed tie-breaking (Baswana et al. 2019; Correa et al. 2019), or dynamically updated priorities (Ehlers et al. 2014).

In these cases, the decision on which type a student should use is often made sequentially, which may not maximize diversity when students have overlapping types. Similar approaches have been employed by Kominers and Sönmez (2013, 2016), Aygün and Bó (2020), and Aygün and Turhan (2020b,a). Sönmez and Yenmez (2022) introduced the concept of "smart reserves" in controlled school choice, offering a more flexible approach to type selection. Aziz and Sun (2021a) expanded on this by proposing a new matching model called matching with multi-rank reserves, providing a comprehensive framework to address various diversity constraints.

Type combinations were previously studied by Aziz, Gaspers, and Sun (2020) under the one-for-all convention, where they derived new quotas for type combinations from original type quotas. However, this method does not fully maximize diversity relative to the original quotas, as it focuses on the new quotas for type combinations, neglecting the original type quotas.

Our new algorithms, based on the ranked reservation graph, build on the concept of rank-maximal matching (Irving et al. 2006). Recent works by Peters (2022) and Hosseini et al. (2021) have applied rank-maximal matching to elicit agent preferences in the house allocation model, which deals with allocating indivisible objects among agents.

Model

In this paper, we focus on how a single school selects students based on its diversity goals while adhering to a priority order, as explored in recent works (Aziz and Sun 2021b; Sönmez and Yenmez 2022; Hafalir et al. 2022; Yokote et al. 2023). In the Appendix, we will also demonstrate how this selection procedure can be seamlessly incorporated into the generalized deferred acceptance algorithm (Hatfield and Milgrom 2005), allowing for the accommodation of multiple schools.

Diversity goals often take the form of minimum and maximum quotas, where schools first aim to fulfill minimum quotas, followed by maximum quotas whenever possible. This approach is known as *dynamic priority* (Ehlers et al. 2014), where precedence is given to students who contribute to achieving diversity goals, even if the resulting set of selected students does not strictly follow the school's priority order over students.

Minimum and maximum quotas can be translated into *two ranks of quotas*, as illustrated in Figure 1. These ranks rep-

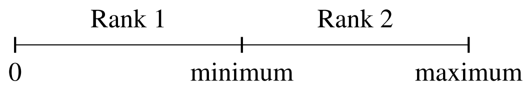

Figure 1: An interpretation of minimum and maximum quotas for type t as two ranks of quotas is as follows. Rank 1 corresponds to the scenario where the number of matched students of type t does not exceed the minimum quota. Rank 2 applies when the number of matched students surpasses the minimum quota but does not exceed the maximum quota.

resent the importance of each quota, with smaller ranks indicating higher importance. Schools prioritize filling quotas with smaller ranks first, consistent with the concept of dynamic priority.

This framework of multiple ranks provides a flexible way to model diverse constraints, including principles such as egalitarianism and proportionality. For further discussion on how such constraints are handled in school choice, we refer readers to (Aziz et al. 2020; Aziz and Sun 2021b).

An instance consists of a tuple $I = (S, q, \succ, T, \eta)$, where S denotes a set of students, q represents the school capacity, $\succ$ is a strict priority order over S, and T denotes a set of types. With a slight abuse of notation, we use $T(s)$ to denote the subset of types associated with student $s \in S$. The parameter η represents a set of *ranked quotas*, where η_t^j denotes the quota for type t at rank j. For simplicity, we consider no more than two ranks of quotas in our examples throughout this paper.

Let $U \subseteq 2^T$ represent the set of type combinations, which are referred to as *groups*. Each student s belongs to a unique group, denoted by $U(s)$. We consider only the groups associated with the students in S, so $|U| \leq |S|$.

Next, we illustrate the issue of imbalance across groups using Example 1. This aspect, overlooked in previous studies, underscores the need for further consideration.

Example 1. Consider a school c with a capacity of 100 seats, which includes minimum quotas of 25 seats each for two types, t_1 and t_2. Additionally, the school reserves 50 seats for a general type t_0, available to any student. We define four groups as follows: u_{00}, u_{10}, u_{01}, and u_{11}, where the subscripts indicate the presence (1) or absence (0) of types t_1 and t_2. Specifically, u_{10} includes type t_1 but not t_2, u_{01} includes type t_2 but not t_1, and u_{11} includes both types, as depicted in Figure 2. Assume there are 50 students in each group, and the priority ordering for students is $u_{00} \succ u_{10} \succ u_{01} \succ u_{11}$. Here, we focus on how many students from each group will be chosen, rather than on which particular students will be selected.

The algorithms by Sönmez and Yenmez (2022) and Aziz and Sun (2021a) produce the same outcome: 50 students from u_{00}, 25 students from u_{10}, 25 students from u_{01}, and 0 students from u_{11} are selected. While the minimum quotas for types t_1 and t_2 are met, the outcome is highly inequitable for students from u_{11}, as none of them are chosen.

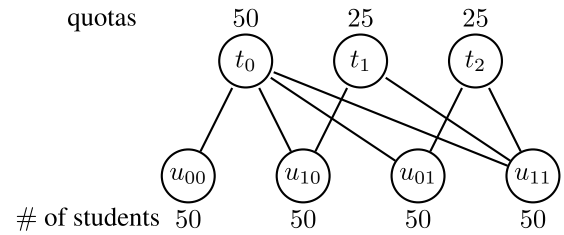

Figure 2: Three types and four groups in Example 1.

Graph Structure

In this section, we introduce two graph structures that underpin our results and algorithms. We begin by revisiting the *ranked reservation graph* from previous work, followed by presenting a more compact representation using a flow network structure.

Ranked Reservation Graph

Irving et al. (2006) studied *ranked bipartite graphs*, where edges are assigned ranks corresponding to students' priorities. Aziz and Sun (2021a) introduced *ranked reservation graphs* to represent diversity goals within the model of multiple ranks of quotas, where the ranks correspond to the quotas rather than to the priorities.

Given an instance $I = (S, q, \succ, T, \eta)$, the corresponding ranked reservation graph $G = (S \cup V, E)$ is a bipartite graph with one set of vertices S representing students and another set V representing reserved seats.

Each reserved seat $v_{t,i}^j$ is associated with a rank j, a type t, and an index i. The indices distinguish between reserved seats of the same rank j and type t. For rank j and type t, there are η_t^j reserved seats. Additionally, a general type t_0 is introduced, available to all students. There are q reserved seats for type t_0 with the largest rank.

An edge $(s, v_{t,i}^j)$ is added between student s and reserved seat $v_{t,i}^j$ if the student has type t. Each edge $(s, v_{t,i}^j)$ is associated with the *rank* j, corresponding to the rank of the reserved seat $v_{t,i}^j$. The set of all edges E can be partitioned into r categories: $E = E^1 \cup \cdots \cup E^r$, where E^j denotes the set of edges with rank j, and r represents the largest rank.

Definition 1 (Matching)**.** Given a graph G, a matching M is a set of edges in which each vertex appears in at most one edge of the matching.

Definition 2 (Signature)**.** Given a ranked reservation graph G, the *signature* of a matching M, denoted as $\rho(M)$, is a tuple of integers $\langle x_1, x_2, \ldots, x_r \rangle$, where each element x_i represents the number of matched edges of rank i in M.

The signatures of two matchings are compared in a lexicographical manner. Specifically, a matching M' with signature $\rho(M') = \langle x_1, \cdots, x_r \rangle$ is *strictly better* than a matching M'' with signature $\rho(M'') = \langle y_1, \cdots, y_r \rangle$ if there exists an index $1 \leq k \leq r$ such that for all indices i where $1 \leq i < k$, $x_i = y_i$ and $x_k > y_k$. A matching M' is *weakly better* than a matching M'' if M'' is not strictly better than M'.

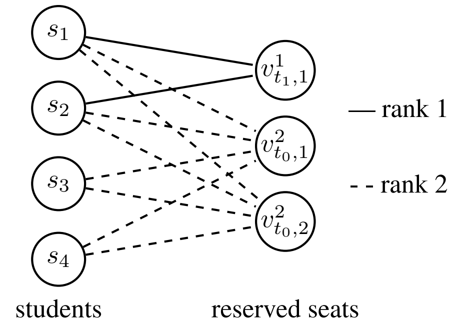

Figure 3: A ranked reservation graph for Example 2. Create one reserved seat $v^1_{t_1,1}$ of rank 1 for type t_1 and two reserved seats $v^2_{t_0,1}$ and $v^2_{t_0,2}$ of rank 2 for type t_0. Solid lines incident to vertex $v^1_{t_1,1}$ have rank 1 and dashed lines incident to vertices $v^2_{t_0,1}$ and $v^2_{t_0,2}$ have rank 2.

Definition 3 (Rank Maximality)**.** Given a ranked reservation graph G, a matching M of size at most q is considered *rank-maximal* if it is weakly better than any other matching of the same size.

The size requirement in the definition of rank maximality is necessary because the total number of reserved seats may exceed the school's capacity when including q reserved seats for the general type t_0.

Example 2. Consider four students $S = \{s_1, s_2, s_3, s_4\}$ and two types: a privilege type t_1 and a general type t_0. Students s_1 and s_2 belong to type t_1, while all students are associated with type t_0. A minimum quota of 1 is imposed on type t_1, and the school has a capacity of 2, with the priority order $s_4 \succ s_3 \succ s_2 \succ s_1$.

The ranked reservation graph, depicted in Figure 3, uses solid lines to represent rank 1 edges and dashed lines for rank 2 edges. The matching $M_1 = \{(s_1, v^1_{t_1,1}), (s_2, v^2_{t_0,1})\}$ has a signature $\rho(M_1) = \langle 1, 1\rangle$ and is rank-maximal. In contrast, the matching $M_2 = \{(s_3, v^2_{t_0,1}), (s_4, v^2_{t_0,2})\}$ has a signature $\rho(M_2) = \langle 0, 2\rangle$.

Flow Network

In this paper, we introduce a more compact graph structure that partitions agents into distinct *groups* based on the types with which each agent is associated. Recall that each distinct type combination is referred to as a group $u \in U$.

Formally, given an instance $I = (S, q, \succ, T, \eta)$, the corresponding flow network $F = (V, E)$ is defined as a directed graph with capacities c and costs (or weights) w. The flow network is organized into four layers, along with two additional vertices: a source vertex s and a sink vertex t.

- *First Layer*: This layer consists of $|U|$ nodes, each corresponding to a group $u \in U$. A directed edge is created from the source vertex s to each node u in this layer, with a capacity equal to $|S_u|$ (the number of students from group u) and a weight of 0.
- *Second Layer*: This layer contains $|T|$ nodes, each representing a type $t \in T$. For every group u that includes type t, a directed edge is added from the node u in the first layer to the node t in the second layer. The capacity of this edge is $|S_t \cap S_u|$ (the number of students with type t from group u), and its weight is 0.
- *Third Layer*: This layer consists of $|T| \times r$ nodes, where each node t^i corresponds to a reserved quota η^i_t with rank i (with the largest rank being r). An edge is drawn from each node t in the second layer to each node t^i in the third layer, with a capacity of η^i_t and a weight of i.
- *Fourth Layer*: This layer includes a single node, denoted as Q, representing the school capacity. Each node t^i in the third layer has a directed edge to the node Q, with a capacity of η^i_t and a weight of 0.
- Finally, a directed edge is added from the node Q to the sink vertex t, with a capacity equal to q (the school capacity) and a weight of 0.

Note that only edges between Layer 2 and Layer 3 have non-zero costs, corresponding to the ranks of the quotas.

Definition 4 (Flow)**.** A *flow* in a network F is a function $f : V \times V \to \mathbb{R}$ that satisfies the following two properties:

1. **Capacity Constraint**: For any $u, v \in V$,
$$0 \leq f(u, v) \leq c(u, v),$$
where $c(u, v)$ denotes the capacity of the edge (u, v).
2. **Flow Conservation**: For all $u \in V \setminus \{s, t\}$,
$$\sum_{v \in V} f(v, u) = \sum_{v \in V} f(u, v).$$

Definition 5 (Minimum Cost Maximum Flow)**.** The *value* of a flow f is defined as
$$|f| = \sum_{v \in V} f(s, v).$$
The *cost* of a flow f is defined as
$$\text{Cost}(f) = \sum_{(v,u) \in E} w(v, u) \times f(v, u).$$
A *minimum cost maximum flow* is a flow that maximizes the flow value while minimizing the total cost among all flows that achieve this maximum value.

Example 3. Figure 4 depicts the flow network of Example 2. Consider the matching $M_1 = \{(s_1, v^1_{t_1,1}), (s_2, v^2_{t_0,1})\}$ with a signature $\rho(M_1) = \langle 1, 1\rangle$ in the ranked reservation graph. It corresponds to a minimum cost maximum flow where two students s_1 and s_2 are sent from s to u_1, student s_1 is then sent through $u_1 \to t_1 \to t^1_1 \to q \to t$ and student s_2 is then sent through $u_1 \to t_2 \to t^2_0 \to q \to t$.

Theorem 1. Given an instance I, let G be the corresponding ranked reservation graph, and F be the associated flow network. A matching M is rank-maximal in G if and only if the corresponding flow f is a minimum-cost maximum flow in F.

Desirable Properties of a Choice Function

In this section, we define some desirable properties of *choice functions* for schools. Intuitively, a choice function selects a

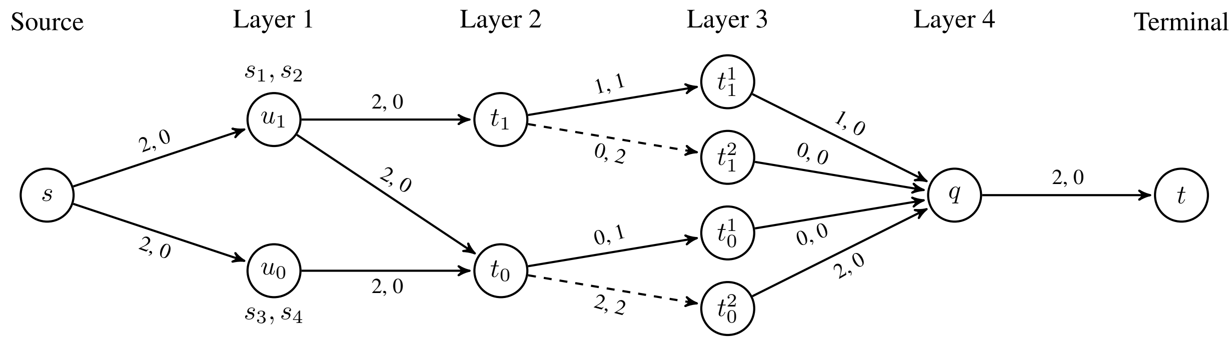

Figure 4: The flow network corresponding to Example 2 consists of four layers, plus a source node and a terminal node. Each edge is labeled with (c, w) where c denotes the edge capacity and w denotes the cost. Only edges between Layer 2 and Layer 3 have non-zero costs, representing their ranks. For example, dashed lines represent a cost of 2.

subset of students from a pool of applicants. For example, in the classical deferred acceptance algorithm (Gale and Shapley 1962), a school's choice function simply selects students based on its priority order, up to its capacity. In this paper, we aim to design a choice function that balances achieving diversity goals with respecting school priorities.

Formally, given an instance $I = (S, q, \succ, T, \eta)$, a choice function $Ch : I \rightarrow 2^S$ returns a subset $S^* \subseteq S$ of students. We assume the choice function satisfies the feasibility requirement: $|Ch(I)| \leq q$.

The first property is a modest efficiency criterion commonly used in two-sided matching models (Goto et al. 2016; Hamada et al. 2017; Kurata et al. 2017; Aziz, Gaspers, and Sun 2020). A choice function is *non-wasteful* if a student is excluded only when the school's capacity q is fully utilized. Here, we assume all applicants are acceptable to the school, or we can exclude all unacceptable students.

Definition 6 (Non-wastefulness)**.** A choice function satisfies non-wastefulness if $|Ch(I)| = \min(|S|, q)$.

Maximal diversity, as introduced by Sönmez and Yenmez (2022), was initially applied to scenarios with a single rank of quotas (i.e., maximum quotas for types). The core idea is that the selected students should be able to form a maximum matching within the reserved seats graph. Aziz and Sun (2021a) extended this concept to models with multiple ranks of quotas by proposing the use of rank-maximal matchings instead of maximum matchings. In essence, a choice function satisfies maximal diversity if it consistently selects a set of students that can be included in a rank-maximal matching.

Definition 7 (Maximal Diversity)**.** A choice function $Ch(I)$ satisfies maximal diversity if, for the corresponding ranked reservation graph G, there exists a rank-maximal matching M of size at most q such that $S_M \subseteq Ch(I)$, where S_M denotes the set of students matched in M.

To promote a more balanced distribution among different groups, we introduce the concepts of *selection ratio* and *balanced representation*. For a given matching M, the selection ratio for group u is defined as the proportion of students selected from group u, i.e., $\frac{|M_u|}{|S_u|}$, where $|M_u|$ is the number of students chosen from group u, and $|S_u|$ is the total number of students in that group.

Intuitively, a choice function achieves balanced representation if it selects a set of students that maximizes the minimum selection ratio across all groups, among the matchings that also satisfy maximal diversity.

Definition 8 (Balanced Representation)**.** A choice function satisfies *balanced representation* if there exists a matching $M \in \mathbb{M}$ such that the set of students selected by the choice function is matched in M and

$$M \in \underset{M' \in \mathbb{M}}{\arg\max} \min_{u \in U} \frac{|M'_u|}{|S_u|}$$

where $\mathbb{M}$ denotes the set of all matchings that adhere to maximal diversity, M'_u represents the subset of students from group u who are matched in M', and S_u denotes the set of all students belonging to group u.

Example 4. Consider Example 1 again. For a balanced representation matching, the minimum selection ratio is 50% for all groups, where the number of matched students from each group is exactly 25.

In the classical school choice problem without diversity constraints (Abdulkadiroğlu and Sönmez 2003), student s is said to have envy towards student s' if student s prefers the school c assigned to s' and has a higher priority at school c than s'. A matching is considered fair if no student has envy towards another. In our setting, we can reinterpret this concept: an unmatched student has envy towards another if the student has a higher priority. However, in general, this fairness concept is incompatible with non-wastefulness and maximal diversity, even if each student is associated with one type (Ehlers et al. 2014).

A natural, less stringent version of fairness restricts envy to occur only within the same group of students.

Definition 9 (Respect of Priorities for Same Groups)**.** A choice function respects priorities for students from the same group if, for any unmatched student s and for any student $s' \in Ch(I)$ with $U(s') = U(s)$, it holds that $s' \succ s$.

We now introduce the concept of justified envy-freeness, which is stronger than respect for priorities within the same

Algorithm 1: Maximum and Balanced Choice Function

Input: An instance $I = (S, q, \succ, T, \eta)$
Output: A set of students $S^* \subseteq S$

```
Selected students S* ← ∅
for each s ∈ S in the descending order of ≻ do
    if there exists a matching M of size at most q that satisfies maximal diversity and balanced representation, and includes all students in S* ∪ {s} then
        S* ← S* ∪ {s}
    end if
end for
return S*
```

group. This concept asserts that a student with higher priority is unmatched only if replacing a student with lower priority would compromise the criteria of maximal diversity or balanced representation.

Definition 10 (Justified Envy-freeness)**.** Given an instance I, a choice function $Ch(I)$ satisfies justified envy-freeness if, for any unselected student $s \notin Ch(I)$, there does not exist a student $s' \in Ch(I)$ such that, if $S' = Ch(I) \cup \{s\} \setminus \{s'\}$, the following conditions hold: (i) $s \succ s'$, and (ii) S' satisfies both maximal diversity and balanced representation.

Although these properties are defined for the choice function of schools, we sometimes use the term loosely by stating that "a matching" satisfies a particular property instead of the choice function. This is because these properties can be framed such that there exists a matching that meets certain conditions in the corresponding ranked reservation graph, with the choice function returning the set of students who are matched in such a matching.

Design of Choice Function

In this section, we first introduce a straightforward choice function that uniquely satisfies the four key properties: maximal diversity, non-wastefulness, justified envy-freeness, and balanced representation. We then provide a high-level overview of how this choice function is implemented. The specific methods for implementation are dependent on the two graph structures we utilize, and we defer the technical details to the subsequent sections.

The new choice function is described in Algorithm 1. The algorithm iterates over each student s based on the school's priority ordering $\succ$ and checks whether s can be matched in a matching of size at most q that achieves maximal diversity and balanced representation along with a set of previously selected students denoted by S^*. If so, s is added to the set of chosen students S^*. Otherwise, the algorithm moves on to the next student. After traversing each student, the algorithm returns S^* as output. The choice function is similar in spirit to (Sönmez and Yenmez 2022) and (Aziz and Sun 2021a).

Example 5. We illustrate Algorithm 1 with Example 2. Recall that school c imposes a minimum quota of 1 for type t_1 and has a capacity of 2. Balanced representation requires that one student is selected from each group. Algorithm 1 checks each student in the decreasing order of school priority: s_4, s_3, s_2, s_1. Student s_4 is first added to S^*, since there exists a matching $M = \{(s_1, v^1_{t_1,1}), (s_4, v^2_{t_0,1})\}$ satisfying maximal diversity and balanced representation. However, adding student s_3 to S^* would violate balanced representation and maximal diversity, so s_3 is not selected. Student s_2 is then added to S^*, since the matching $M = \{(s_2, v^1_{t_1,1}), (s_4, v^2_{t_0,1})\}$ satisfies maximal diversity and balanced representation. Student s_1 cannot be selected due to the capacity restriction. Thus, Algorithm 1 returns $S^* = \{s_2, s_4\}$ as the outcome.

We next demonstrate that the choice function we propose is the unique one satisfying all four desirable properties.

Theorem 2. The choice function in Algorithm 1 is the unique one which satisfies non-wastefulness, maximal diversity, justified envy-freeness and balanced representation.

High Level Description of Implementation

Next, we introduce two concepts that are critical for the effective implementation of Algorithm 1. Intuitively, an instance is considered *valid* with respect to a target vector for groups if it allows for a matching that maximizes diversity while ensuring that the number of matched students from each group meets the specified targets.

Definition 11 (Validity)**.** An instance $I = (S, q, \succ, \eta)$ is considered *valid* with respect to a target vector $\delta = (\delta_u)_{u \in U}$ if there exists a matching M such that:
1. M achieves maximal diversity, and
2. For all $u \in U$, $|M \cap S_u| \geq \delta_u$, where S_u denotes the set of students from group u.

If we have an algorithm Γ that can check in polynomial time whether an instance is valid with respect to a given target vector for groups, we can use binary search, invoking Γ, to compute the max-min selection ratio needed to achieve balanced representation, as shown shortly.

Definition 12 (Max-min Ratio & Crucial Vector)**.** Given an instance I, let α denote the max-min selection ratio that ensures balanced representation. Let $\delta^* = (\delta_u)_{u \in U}$ be the crucial vector corresponding to the max-min selection ratio α, where $\delta_u = \lfloor \alpha \cdot |S_u| \rfloor$ for all $u \in U$.

The following Theorem 3 asserts that if an instance I is valid with respect to the crucial vector (corresponding to balanced representation), then a matching exists that satisfies all four properties. This theorem provides a sufficient condition to guarantee the existence of such a matching.

Theorem 3. An instance I is valid with respect to its crucial vector δ^* if and only if it admits a matching that satisfies non-wastefulness, maximal diversity, balanced representation, and justified envy-freeness.

We now present a high-level overview of the implementation of the choice function in Algorithm 1 by addressing the following two key questions:

- The **first question** involves designing an efficient algorithm to check whether an instance is valid, which allows us to compute the max-min selection ratio among all groups using binary search.

Algorithm 2: Checking Whether an Instance is Valid

Input: An instance $I = (S, q, \succ, T, \eta)$ and a target vector $\delta = (\delta_u)_{u \in U}$

Output: a rank-maximal matching M of size at most q s.t. $\forall u \in U, |M_u| \geq \delta_u$ or no-instance

1: Construct a ranked reservation graph G and compute a rank-maximal matching M of size at most q
2: Denote the set of students matched in M as S_M
3: $\widetilde{S} \leftarrow \emptyset$
4: **for** each group $u \in U$ **do**
5:     Choose $\min(\delta_u, |S_{M,u}|)$ students from $S_{M,u}$ (the set of students from group u in S_M) and add them to $\widetilde{S}$
6: **end for**
7: **for** each group $u \in U$ **do**
8:     **while** $|\widetilde{S}_u| < \delta_u$ **and** $\exists s \in S_u \setminus S_M$ **do**
9:         **if** there exists an *alternating path* P w.r.t. M that starts from s and ends at $s' \in S_M \setminus \widetilde{S}$ **then**
10:             $\widetilde{S} \leftarrow \widetilde{S} \cup \{s\}, M \leftarrow M \oplus P$
11:         **else if** there exist i) an *augmenting path* P w.r.t. M that starts from s and ends at some unmatched seat v of rank k and ii) a student $s' \in S_M \setminus \widetilde{S}$ who is matched to some seat v' of rank k in M **then**
12:             $\widetilde{S} \leftarrow \widetilde{S} \cup \{s\}, M \leftarrow M \oplus P, M \leftarrow M \setminus \{(s', v')\}$
13:         **end if**
14:     **end while**
15:     **if** $|\widetilde{S}_{M,u}| < \delta_u$ **then**
16:         **return** NO-instance
17:     **end if**
18: **end for**
19: **return** matching M

- The **second question** focuses on designing an efficient algorithm to verify the existence of a matching with the desired properties, while ensuring that a specific set of students is included (line 3). We can then repeat this procedure and update the matching accordingly to ensure justified envy-freeness, while maintaining both maximal diversity and balanced representation.

Algorithm Design Based on Rank-maximal Matching

In this section, we propose two polynomial-time algorithms based on rank-maximal matching to solve the two technical questions in implementing Algorithm 1.

Given a matching M, an *alternating path* is a set of edges that begins with an unmatched vertex and whose edges belong alternately to the matching and not to the matching. An *augmenting path* is an alternating path that starts and ends on unmatched vertices. Given an augmenting path P, let $M \oplus P$ denote a new matching in which edges from $P \setminus M$ are matched while edges from $M \cap P$ are not matched.

Checking Validity

Algorithm 2 solves the first question of checking whether a given instance is valid w.r.t. some target vector δ and it returns a rank-maximal matching for a yes-instance.

Theorem 4. Given an instance I and a target vector δ, Algorithm 2 checks whether the instance I is valid w.r.t. δ in time $O(r|E|\sqrt{|V|} + |E||V|)$, where r, V and E denote the number of ranks, the number of nodes and the number of edges in the ranked reservation graph.

Example 6. We illustrate how Algorithm 2 works through Example 2. Suppose school c imposes a minimum target of 1 for both groups u_1 and u_0. First, we compute a rank-maximal matching of size at most 2, say $M = \{(s_1, v^2_{t_0,1}), (s_2, v^1_{t_1,1})\}$. Both students s_1 and s_2 are from group u_1, and we need to select one of them. Let's choose s_2 from $S_M = \{s_1, s_2\}$ (where S_M is the set of students matched in M), so we set $\widetilde{S} = \{s_2\}$. For the undersubscribed group u_0, we select one unmatched student, say s_3, and check whether s_3 can be added to $\widetilde{S}$. Since there exists an alternating path from s_3 to s_1 through the seat $v^2_{t_0,1}$, we can add s_3 to $\widetilde{S}$. Consequently, we update the matching to $M' = \{(s_3, v^2_{t_0,1}), (s_2, v^1_{t_1,1})\}$. Thus, Example 2 is valid and Algorithm 2 returns matching M' as outcome.

Implementation Based on Rank-maximal Matching

We next provide a detailed implementation of the choice function in Algorithm 1, based on a rank reservation graph, as described in Algorithm 3. The input consists of an instance I, a critical vector δ^*, and a rank-maximal matching M of size at most q that is valid w.r.t. δ^* and the top δ^*_u students are matched for each $u \in U$.

Example 7. We illustrate how Algorithm 3 works using Example 2. Suppose we start with the rank-maximal matching $M = \{(s_3, v^2_{t_0,1}), (s_2, v^1_{t_1,1})\}$, which was produced by Algorithm 2 as described in Example 6. We first update M to ensure that the top δ^*_u students are matched, resulting in $M = \{(s_2, v^1_{t_1,1}), (s_4, v^2_{t_0,1})\}$. We then check all students in the order s_4, s_3, s_2, s_1. Student s_4 is already matched and is added to S^*. Student s_3 cannot be added, as none of the required conditions are met. Student s_2 is already matched in M and is added to S^*. Student s_1 is rejected, and the algorithm terminates.

Algorithm Design Based on Flow Network

In this section, we propose new algorithms to implement the choice function in Algorithm 1 using flow networks. Compared to the methods based on a ranked reservation graph, these algorithms are more straightforward and much faster for large instances where the number of students significantly exceeds the number of groups.

To distinguish it from the ranked reservation graph G, we denote the number of nodes and edges in the flow network F as n and m, respectively. Specifically, $n = O(|T|r + |U|)$ and $m = O(|U||T| + r|T|^2)$, where $|U| \leq 2^{|T|} \leq |S|$.

Theorem 5. Given an instance I and a target vector δ, let F denote the corresponding flow network. Checking validity with respect to δ can then be done in time $O(m \log(n)(m + n \log(n)))$, where m and n denote the number of edges and nodes in the flow network F.

Algorithm 3: Implementation of the Choice Function in Algorithm 1 Based on Rank-maximal Matching

Input: An instance $I = (S, q, \succ, T, \eta)$, a critical vector δ^*, a rank-maximal matching M of size at most q that is valid w.r.t. δ^*, and the top δ^*_u students are matched.
Output: a set of students S^*

Selected students $S^* \leftarrow \emptyset$
for each $s \in S$ in the descending order of $\succ$ **do**
  if $s \in S_M$ **then**
    $S^* \leftarrow S^* \cup \{s\}$
  else
    for $s' \in S_M \cap S_u \setminus S^*$ s.t. $|S_M \cap S_u| > \delta^*_u$ **do**
      if there exists an *alternating path* P w.r.t. M that starts from s and ends at s' **then**
        $S^* \leftarrow S^* \cup \{s\}$, $M \leftarrow M \oplus P$
        break
      else if i) there exists an *augmenting path* P w.r.t. M that starts from s and ends at some unmatched seat v of rank k and ii) student s' is matched to some seat v' of rank k in M **then**
        $S^* \leftarrow S^* \cup \{s\}$
        $M \leftarrow M \oplus P$, $M \leftarrow M \setminus \{(s', v')\}$
        break
      end if
    end for
  end if
end for
return a set of students S^*

Computing a Crucial Vector

We next proceed to compute a crucial vector in Algorithm 4.

Theorem 6. Given an instance I, Algorithm 4 computes a crucial vector in time $O(m \log(n)(m + n \log(n))) \log(|S|)$ where m and n denote the number of edges and nodes in the flow network.

Implementation Based on Flow Network

We next proceed to the implementation of the choice function in Algorithm 1 using a flow network approach. The algorithm first selects the top δ^*_u students from each group S_u based on the priority order $\succ$. For each unselected student $s \in S \setminus S^*$, we then check whether s can be included while still satisfying maximal diversity and balanced representation. This is done by verifying if there exists a minimum cost maximum flow with a lower bound of $|S'_u|$ on each edge from the source node to the node u, where $S' = S^* \cup \{s\}$. If such a flow exists, s can be included in the selection.

Theorem 7. Given a crucial vector δ^*, Algorithm 5 returns a matching satisfying maximal diversity, non-wastefulness, balanced representation and justified envy-freeness in time $O(|S| m \log(n)(m + n \log(n)))$ where m and n denote the number of edges and nodes in the flow network.

If we assume that the number of types and type combinations is bounded by a small constant, as is often the case in practical markets, then solving the problem based on the flow network structure is much faster than solving it using the ranked reservation graph.

Algorithm 4: Computing a Crucial Vector

Input: An instance $I = (S, q, \succ, T, \eta)$
Output: A crucial vector $\delta^* = (\delta^*_u)_{u \in U}$

Initialize $\delta^*_u \leftarrow 0$ for each group $u \in U$
Initialize $\alpha_0 \leftarrow 0$ and $\alpha_1 \leftarrow 1$
while some δ^*_u differs from the last round **do**
  $\alpha \leftarrow (\alpha_0 + \alpha_1)/2$
  $\delta^*_u \leftarrow \lfloor \alpha \cdot |S_u| \rfloor, \forall u \in U$
  Construct a flow network F. For each edge from s to u in the first layer, add a lower bound δ^*_u
  if there exists a minimum cost maximum flow **then**
    $\alpha_0 \leftarrow \alpha$ {Search between α and α_1}
  else
    $\alpha_1 \leftarrow \alpha$ {Search between α_0 and α}
  end if
end while
return δ^*

Algorithm 5: Implementation of the Choice Function in Algorithm 1 Based on Flow Network

Input: An instance $I = (S, q, \succ, T, \eta)$ and a crucial vector $\delta^* = (\delta_u)_{u \in U}$
Output: A set of students S^*

Initialize the selected students, S^*, as the union of the top δ_u students from each group S_u according to $\succ$ for all $u \in U$.
Construct a flow network F
for each $s \in S \setminus S^*$ in descending order of $\succ$ **do**
  Set $S' \leftarrow S^* \cup \{s\}$
  For each edge from s to u in the first layer, add a lower bound $|S'_u|$
  if a minimum cost maximum flow exists **then**
    Update $S^* \leftarrow S^* \cup \{s\}$
  end if
end for
return S^*

Conclusion

In this paper, we have addressed the challenge of student placements under diversity constraints by introducing the concept of *balanced representation*. Our proposed choice function is distinguished by its satisfaction of maximal diversity, non-wastefulness, balanced representation, and justified envy-freeness. We demonstrated the practical advantages of our approach through the development of efficient algorithms. Future work could explore further refinements to these algorithms and investigate their applicability to other settings where fairness and efficiency are critical.

The complete version of this paper is accessible at https://arxiv.org/abs/2412.13622.

Acknowledgments

This work was partially supported by the Japan Science and Technology Agency under the ERATO Grant Number JPMJER2301. We thank the anonymous reviewers of AAAI 2025 for their valuable comments.

References

Abdulkadiroğlu, A. 2005. College admissions with affirmative action. *International Journal of Game Theory*, 33(4): 535–549.

Abdulkadiroğlu, A.; and Sönmez, T. 2003. School Choice: A Mechanism Design Approach. *American Economic Review*, 93(3): 729–747.

Aygün, O.; and Bó, I. 2020. College admission with multidimensional privileges: The Brazilian affirmative action case. *American Economic Journal: Microeconomics*.

Aygün, O.; and Turhan, B. 2016. Dynamic reserves in matching markets: Theory and applications.

Aygün, O.; and Turhan, B. 2017. Large-Scale Affirmative Action in School Choice: Admissions to IITs in India. *American Economic Review*, 107(5): 210–13.

Aygün, O.; and Turhan, B. 2020a. Designing Direct Matching Mechanism for India with Comprehensive Affirmative Action. Papers 2004.13264, arXiv.org.

Aygün, O.; and Turhan, B. 2020b. Dynamic reserves in matching markets: Theory and applications. *Journal of Economic Theory*, 188.

Aziz, H. 2019. A Rule for Committee Selection with Soft Diversity Constraints. *Group Decision and Negotiation*, 1—8.

Aziz, H.; and Brandl, F. 2021. Efficient, Fair, and Incentive-Compatible Healthcare Rationing. In *Proceedings of the 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, EC 2021*, 103–104. ACM.

Aziz, H.; Gaspers, S.; and Sun, Z. 2020. Mechanism Design for School Choice with Soft Diversity Constraints. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, 153–159.

Aziz, H.; Gaspers, S.; Sun, Z.; and Walsh, T. 2019. From Matching with Diversity Constraints to Matching with Regional Quotas. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 377–385.

Aziz, H.; Gaspers, S.; Sun, Z.; and Yokoo, M. 2020. Multiple Levels of Importance in Matching with Distributional Constraints. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS), Extended Abstract*, 1759–1761.

Aziz, H.; and Sun, Z. 2021a. Multi-rank Smart Reserves. In *Proceedings of the 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, EC 2021*, 105–124.

Aziz, H.; and Sun, Z. 2021b. School Choice with Flexible Diversity Goals and Specialized Seats. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, 17–23. ijcai.org.

Aziz, H.; and Sun, Z. 2025. Multi-rank smart reserves: A general framework for selection and matching diversity goals. *Artificial Intelligence*, 339: 104274.

Baswana, S.; Chakrabarti, P. P.; Chandran, S.; Kanoria, Y.; and Patange, U. 2019. Centralized Admissions for Engineering Colleges in India. In *Proceedings of the 20th ACM Conference on Economics and Computation*, 323–324.

Biró, P.; Fleiner, T.; Irving, R. W.; and Manlove, D. F. 2010. The College Admissions problem with lower and common quotas. *Theoretical Computer Science*, 411(34): 3136 – 3153.

Correa, J.; Epstein, N.; Epstein, R.; Escobar, J.; Rios, I.; Aramayo, N.; Bahamondes, B.; Bonet, C.; Castillo, M.; Cristi, A.; Epstein, B.; and Subiabre, F. 2022. School choice in Chile. *Operations Research*, 70(2): 1066–1087.

Correa, J. R.; Epstein, R.; Escobar, J.; Rios, I.; Bahamondes, B.; Bonet, C.; Epstein, N.; Aramayo, N.; Castillo, M.; Cristi, A.; and Epstein, B. 2019. School Choice in Chile. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019*, 325–343.

Echenique, F.; and Yenmez, M. B. 2015. How to Control Controlled School Choice. *American Economic Review*, 105(8): 2679–94.

Ehlers, L.; Hafalir, I. E.; Yenmez, M. B.; and Yildirim, M. A. 2014. School choice with controlled choice constraints: Hard bounds versus soft bounds. *Journal of Economic Theory*, 153: 648–683.

Gale, D.; and Shapley, L. S. 1962. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1): 9–15.

Gonczarowski, Y. A.; Nisan, N.; Kovalio, L.; and Romm, A. 2019. Matching for the Israeli "Mechinot" Gap Year: Handling Rich Diversity Requirements. In *Proceedings of the 20th ACM Conference on Economics and Computation*, 321–321.

Goto, M.; Iwasaki, A.; Kawasaki, Y.; Kurata, R.; Yasuda, Y.; and Yokoo, M. 2016. Strategyproof matching with regional minimum and maximum quotas. *Artificial Intelligence*, 235: 40–57.

Hafalir, I.; Kojima, F.; Yenmez, B.; and Yokote, K. 2022. Design on Matroids: Diversity vs. Meritocracy. *arXiv preprint arXiv:2301.00237*.

Hafalir, I. E.; Yenmez, M. B.; and Yildirim, M. 2013. Effective affirmative action in school choice. *Theoretical Economics*, 8(2): 325–363.

Hamada, N.; Hsu, C.; Kurata, R.; Suzuki, T.; Ueda, S.; and Yokoo, M. 2017. Strategy-proof school choice mechanisms with minimum quotas and initial endowments. *Artificial Intelligence*, 249: 47–71.

Hatfield, J. W.; and Milgrom, P. R. 2005. Matching with contracts. *American Economic Review*, 95(4): 913–935.

Hosseini, H.; Menon, V.; Shah, N.; and Sikdar, S. 2021. Necessarily Optimal One-Sided Matchings. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, Virtual Event, February 2-9, 2021*, 5481–5488. AAAI Press.

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

Kamada, Y.; and Kojima, F. 2018. Stability and strategy-proofness for matching with constraints: A necessary and sufficient condition. *Theoretical Economics*, 13(2): 761–793.

Kojima, F. 2012. School choice: Impossibilities for affirmative action. *Games and Economic Behavior*, 75(2): 685–693.

Kominers, S. D.; and Sönmez, T. 2013. Designing for diversity in matching. In *Proceedings of the fourteenth ACM Conference on Electronic Commerce, EC 2013, Philadelphia, PA, USA, June 16-20, 2013*, 603–604.

Kominers, S. D.; and Sönmez, T. 2016. Matching with slot-specific priorities: Theory. *Theoretical Economics*, 11(2): 683–710.

Kurata, R.; Hamada, N.; Iwasaki, A.; and Yokoo, M. 2015. Controlled school choice with soft bounds and overlapping types. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*, 951–957.

Kurata, R.; Hamada, N.; Iwasaki, A.; and Yokoo, M. 2017. Controlled school choice with soft bounds and overlapping types. *Journal of Artificial Intelligence Research*, 58: 153–184.

Pathak, P. A.; Sönmez, T.; Ünver, M. U.; and Yenmez, M. B. 2021. Fair Allocation of Vaccines, Ventilators and Antiviral Treatments: Leaving No Ethical Value Behind in Health Care Rationing. In *Proceedings of the 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, EC 2021*, 785–786. ACM.

Peters, J. 2022. Online Elicitation of Necessarily Optimal Matchings. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, Virtual Event, February 22 - March 1, 2022*, 5164–5172. AAAI Press.

Sönmez, T.; and Yenmez, M. B. 2022. Affirmative Action in India via Vertical and Horizontal Reservations. *Econometrica*, 90(3): 1143–1176.

Yokote, K.; Hafalir, I.; Kojima, F.; and Yenmez, B. 2023. Representation Theorems for Path-Independent Choice Rules. *arXiv preprint arXiv:2303.00892*.