

School Choice with Flexible Diversity Goals and Specialized Seats

Haris Aziz¹, Zhaohong Sun¹

¹UNSW Sydney

{haris.aziz, zhaohong.sun}@unsw.edu.au

Abstract

We present a new and rich model of school choice with flexible diversity goals and specialized seats. The model also applies to other settings such as public housing allocation with diversity objectives. Our method of expressing flexible diversity goals is also applicable to other settings in moral multi-agent decision making where competing policies need to be balanced when allocating scarce resources. For our matching model, we present a polynomial-time algorithm that satisfies desirable properties, including strategyproofness and stability under several natural subdomains of our problem. We complement the results by providing a clear understanding about what results do not extend when considering the general model.

1 Introduction

Diversity goals are prevalent in many scenarios including the hiring of employees, student-intake, and public housing [Benabbou *et al.*, 2018]. These may be affirmative action legal requirements, institutional policies, or guidelines for ensuring a healthier and better balance of various groups. As important decisions are made concerning who gets the next job or who gets their preferred school seat, it is critical that the algorithms are transparent and fair. This is especially so when they make complex decisions that simultaneously take into account merit, priority, and diversity.

In the past few years, diversity issues have been heavily studied in school choice [Abdulkadiroğlu and Sönmez, 2003], the prototypical model of two-sided matching. In the basic school choice problem, students have preferences over schools and schools have priorities over students. The typical goal is to match the students to schools in a way that is stable (no student wants to take an vacant slot or wants to replace a lower priority student). Another property that is important in this context is strategyproofness (no student has an incentive to misreport her preferences).

In school choice *with diversity goals*, schools accept students while not just taking into account the priority ordering over students (that could be based on test scores) but also considering diversity requirements (such as keeping in mind

a target number of students with special talents or disadvantaged backgrounds). Student placement under diversity concerns has been considered in many countries including Israel [Gonczarowski *et al.*, 2019] and India [Baswana *et al.*, 2019]. A prominent and successful approach for these problems is to specify diversity requirements via upper and lower bounds on the types of students at the school (see, e.g. Ehlers *et al.* [2014] and Aziz *et al.* [2019]). Based on these quotas, diversity is achieved by first admitting students whose types have not reached the minimum quotas and then admitting students whose types have not reached the maximum quotas [Kurata *et al.*, 2017; Aziz *et al.*, 2020; Sun, 2019].

There are limitations of specifying diversity goals just by imposing upper and lower bounds on types. The first concern is whether one should indeed treat the status of being under-subscribed equally. For instance, suppose one school imposes minimum quota 10 on two types. Type t_1 admits just 1 student while type t_2 admits 8 students. It is reasonable to give higher precedence to students of type t_1 , who will help to achieve a diversity balance. The second concern is whether we should treat all types equally. It is common that some types are more important than others. Thirdly, lower and upper quotas are not sufficient to target proportionality ratios of types (that captures desirable distributions of different types of students). In particular, having targets on absolute numbers does not achieve proportionality goals effectively.

We explain the third issue through the following example. Suppose there is a market where students are associated with one of the three types t_1, t_2 , and t_3 . Consider one school c with capacity 100 that has the target proportions of students as follows: $t_1 : 30\%, t_2 : 30\%, t_3 : 40\%$. If we wish to capture these targets via setting of lower bounds, one can impose minimum quotas 30, 30 and 40 on type t_1, t_2, t_3 respectively. Suppose 15, 60, and 60 students of type t_1, t_2, t_3 apply for the school c . Then notable algorithms for the problem that use minimum quotas would return an outcome with type distribution $(t_1 : 15, t_2 : 45, t_3 : 40)$. However, there exists an outcome with type distribution $(t_1 : 15, t_2 : 37, t_3 : 48)$ that is closer to the target proportions. Such an outcome will be computed by the algorithm we propose. In this paper, we explore the issue of capturing a rich class of diversity goals in matching problems.

In this paper, we aim to address the following questions. *What is a powerful yet computationally tractable framework*

of specifying diversity goals, that encompasses most of the particular diversity approaches in theory and practice? Can it also capture meaningful and well-motivated diversity goals not yet addressed in the literature? Can we use such a framework and achieve desirable objectives in two-sided matching market design such as stability and strategyproofness in a robust manner? What are the limits of such an approach when students are allowed to express preferences over individual school seats?

Contributions We present a new model of matching markets that generalizes school choice with diversity goals in two important ways. (1) Firstly, we allow schools to have specialized *seats* that are motivated by additional features such as accompanying scholarships, tuition-fee waiver etc. (2) Secondly, we greatly expand the type of diversity goals that can be specified. Our model introduces a flexible and general ordinal approach to specifying diversity goals that capture lower and upper quotas as well as proportionality concerns. We discuss how existing approaches that only minimum and maximum quotas in the intended rations may not achieve outcomes that are closest to the intended ratios. *The approach applies to many other settings including social choice and multi-agent decision making in which diversity is a concern.*

For the new model, we present an algorithm called *Generalized Deferred Acceptance with Flexible Diversity (GDA-FD)*. It is a novel algorithm that we design for schools to choose students while taking into account flexible diversity goals. GDA-FD satisfies a natural notion of stability, strategyproofness, and non-wastefulness when each agent has one type and there are no specialized seats.

We then show that allowing for specialized seats but enforcing diversity over students leads to several challenges in achieving stability and strategyproofness. In view of these challenges, we then turn to an important domain in which we allow for specialized seats but there is exactly one school.¹ In this domain, we show that even if students have overlapping types, GDA-FD algorithm satisfies weak Pareto optimality and two notions of strategyproofness. En route to proving these results, we provide an alternative view of GDA-FD by showing its equivalence with an algorithm called *Sequential Allocation Under Dynamic Diversity Goals*.

Finally, we turn to a class of preferences called *school-based* that generalizes the case of homogenous school seats as well as the case of exactly one school. For school-based preferences, we show that GDA-FD is not strategyproof even if each agent has a single type. However, strategyproofness is regained by using another algorithm called *Modified GDA-FD* that calls GDA-FD on a modification of the problem instance. Since Modified GDA-FD is equivalent to GDA-FD when there are no special seats, our key message is that the modified GDA-FD works well on several important domains. Some of our results are summarized in Table 1.

2 Related Work

Abdulkadiroğlu and Sönmez [2003] presented one of the seminal works on the use of matching market design for

¹The domain captures various committee selecting or hiring scenarios that take into account diversity goals.

Preference Domain	Strategyproofness	Stability	Non-wastefulness
No Specialized Seats	✓	✓	✓
One School	✓	✗	✓
School-based	✓	✗	✓

Table 1: Properties satisfied by Modified GDA-FD (Algorithm 4) under different preference domains when each agent has a single type. The positive results for the case of one school also hold for multiple types.

school choice. The problem of school choice with diversity constraints or goals is an active area of research in market design. Typically, each school imposes a maximum quota and a minimum quota on each type [Kojima, 2012; Hafalir *et al.*, 2013; Kominers and Sönmez, 2013]. Since stable matchings cannot exist for hard lower quotas and the corresponding problems become NP-hard (see e.g., Kurata *et al.* [2017] and Chen *et al.* [2020]), most of the successful approaches treat the quotas as soft (see, e.g. Ehlers *et al.* [2014] and followup works). Most work on school choice with diversity goals focuses on minimum and maximum quotas only [Ehlers *et al.*, 2014; Gonczarowski *et al.*, 2019; Sönmez and Yenmez, 2019; Baswana *et al.*, 2019; Aziz *et al.*, 2019]. Although there is some recent work on ratio constraints [Nguyen and Vohra, 2019], our diversity goals are more general and flexible, and our algorithmic solutions are different as well. Kominers and Sönmez [2016] considered a matching model in which each school seat may have a different priority ordering over students and students have rankings over schools. In contrast, we allow students to distinguish between seats of the same school. One more difference from the work of Kominers and Sönmez [2016] is that we take different forms of diversity goals into consideration.

Our model also bears some similarities with hospital-resident matching with regional constraints [Kamada and Kojima, 2015, 2017], in which students are viewed as doctors, school seats are viewed as hospitals and schools are viewed as regions. However, in the hospital-resident matching, the distributional constraints are imposed on the number of doctors matched to different regions but not on the proportional composition of types of doctors. Ahmadi *et al.* [2020] considered diversity goals in a different model that does not involve preferences and priorities.

3 A Model for School Choice with Flexible Diversity Goals and Specialized Seats

In this section, we formalize the model of School Choice with Flexible Diversity Goals and Specialized Seats. An instance I of the setting consists of a tuple $(S, T, H, C, \mathcal{X}, \succ_S, \succ_C, R_C)$. There is a set of students denoted by $S = \{1, \dots, n\}$, each of whom belongs to some of the types from the type space $T = \{t_1, \dots, t_k\}$. Let $T(i)$ denote the set of types of student i . We will denote by C the set of schools with a generic school denoted by b_j or b . We will denote by H the set of school seats with a generic seat denoted by h . School seats H are partitioned into $|C|$ schools, where $b_j \subseteq H$ is the set of school seats controlled by school $b_j \in C$. For

any two different schools $b_i, b_j \in C$, we have $b_i \cap b_j = \emptyset$. The symbol R_C denotes the precedence profile of schools that captures flexible diversity goals. Section 4 is devoted to the precedence profile R_C .

We follow the model of matching with contracts [Hatfield and Milgrom, 2005] where each contract $x = (i, h, b)$ is a student–seat–school tuple indicating that student i is matched to seat h at school b . The set of all contracts is denoted by $\mathcal{X} \subseteq S \times H \times C$ and we assume that for any $(i, h, b) \in \mathcal{X}$, we have $h \in b$. Given any $X \subseteq \mathcal{X}$, for each student i , we denote by X_i as the set of contracts pertaining to student i ; for each seat h , we denote by X_h as the set of contracts pertaining to seat h ; for each school b , we denote by X_b as the set of contracts pertaining to school b . We denote by $X_b^t = \{(i, h, b) \in X_b \mid t \in T(i)\}$ as the set of contracts involving students of type t and school b . We extend the notation of each individual to a set of agents by taking the union, i.e., for any $S' \subseteq S$, $X_{S'} = \bigcup_{i \in S'} X_i$.

The preference profile of students S is denoted by $\succ_S = \{\succ_1, \dots, \succ_n\}$ where for student $i \in S$, each \succ_i specifies the preference of student i over contracts \mathcal{X}_i . For instance, $(i, h_1, b_1) \succ_i (i, h_2, b_2)$ means that student i weakly prefers the contract (i, h_1, b_1) to the contract (i, h_2, b_2) . Let \succ_i and \sim_i denote the strict and indifferent relation, respectively. In Section 5.2, we will assume that a student is indifferent between all contracts involving seats of the same school. In general, if some algorithm of ours requires strict preferences, we will assume that the ties are broken arbitrarily.

The priority profile of schools C is denoted by $\succ_C = \{\succ_{b_1}, \dots, \succ_{b_{|C|}}\}$ where each \succ_b specifies the priority ordering of school b over contracts \mathcal{X}_b . We will assume that a school's priority over contracts is based on an underlying priority relation over students, which could be based on first-come first-served basis, entrance exam scores or randomization.

Next, we introduce several important properties for a desirable outcome. A contract (i, h, b) is *acceptable* to student i and school b if both $(i, h, b) \succ_i \emptyset$ and $(i, h, b) \succ_b \emptyset$ hold where \emptyset represents the option of being unmatched. W.L.O.G, we consider acceptable contracts only, since we can remove any unacceptable contract from \mathcal{X} . An *outcome* or a *matching* X is a set of contracts, i.e. $X \subseteq \mathcal{X}$. An outcome is *feasible* if each student is matched to at most one contract, i.e. $\forall i \in S, |X_i| \leq 1$, and each seat is matched to at most one student, i.e. $\forall h \in H, |X_h| \leq 1$. A feasible outcome Y is *non-wasteful* if there exists no $(i, h', b') \in Y$ such that there exists another contract (i, h, b) with $(i, h, b) \succ_i (i, h', b')$ and $Y \cup \{(i, h, b)\} \setminus \{(i, h', b')\}$ is feasible. Given an instance I , an algorithm is *strategyproof* for students if no student can be matched to a strictly better contract when the student misreports his preferences.

4 Framework for Flexible Diversity Goals

In this section, we present a novel way for schools to select students. A key component of the method is to design a novel method to specify dynamic priorities of a school over students. The dynamic priorities of a school b are based on the static priority relation \succ_b over students, and a dynamic precedence relation R_b over the types in T that we specify below.

Given a feasible outcome X , each school b has a *dynamic precedence ordering* $R_b(X)$ over types T . If $t_1 R_b(X) t_2$ holds, then for the outcome X , type t_1 is weakly preferred to type t_2 in terms of achieving diversity goals of school b . Let P_b and I_b denote the strict and indifferent relation, respectively. We denote by $R_C = \{R_{b_1}, \dots, R_{b_{|C|}}\}$ the precedence profile of all schools.

4.1 Dynamic Precedence Representation

Next, we explain how we represent the dynamic precedence R_b over types. First, each school b specifies different levels over each type t , denoted by $L_{b,t}$ where each level $\ell_{b,t}^j \in L_{b,t}$ specifies the range of number of students of type t that are matched to school b . For instance, in Figure 1, $\ell_{b,t}^1 = [0, 10)$ means that when the number of students of type t matched to school b is weakly larger than 0 and smaller than 10, then it falls into the level $\ell_{b,t}^1$. We refer to $L_b^j = \{\ell_{b,t}^j\}_{t \in T}$ as level j of school b for convenience. Note that for each type t , it is not necessary to define levels for every integer from $[1, \dots, |S|]$.




Figure 1: An instance of levels of type t at school b .

Then the dynamic precedence ordering $R_b(X)$ is determined as follows with indifference classes in decreasing order of precedence from left to right: $R_b(X) : \{t \in T : |X_b^t| \in \ell_{b,t}^1\}, \{t \in T : |X_b^t| \in \ell_{b,t}^2\}, \dots, \{t \in T : |X_b^t| \in \ell_{b,t}^k\}$, where $|X_b^t|$ denotes the number of students of type t matched to school b in the outcome X . Intuitively, for a given outcome X , school b gives the highest precedence to types that fall into level 1 of school b , the second highest precedence to types that fall into level 2 of school b and so on.

Next, we explain how to capture different diversity goals through the dynamic precedence relation $R_b(X)$.

Lexicographic: Let there be a fixed precedence ordering over types, say $t_1, \dots, t_{|T|}$. The lexicographic diversity goal of school b is that whenever possible, it prefers to be matched to students of type t with smaller index. In that case, for each type t_i , create one level $\ell_{b,t_i}^i = [0, |S|]$.

Min and Maximum quotas: For each school b , let η_b^t and $\bar{\eta}_b^t$ denote the minimum quota and the maximum quota of type t , respectively. The diversity goal of school b under minimum and maximum quotas is that, school b gives the highest precedence to the student whose type has not reached the minimum quota, medium precedence to the student whose type has reached the minimum quota but not the maximum quota, and the lowest precedence to the student whose type has reached the maximum quota [Ehlers *et al.*, 2014]. In that case, for each type t , school b has three levels $\ell_{b,t}^1 = [0, \eta_b^t - 1]$, $\ell_{b,t}^2 = [\eta_b^t, \bar{\eta}_b^t - 1]$ and $\ell_{b,t}^3 = [\bar{\eta}_b^t, |S|]$. When only minimum quotas or maximum quotas exist, we can reduce to two levels.

Proportional Let there be $|T|$ positive integers, say $r_1, \dots, r_{|T|}$ corresponding to each type t . The proportional diversity goal of school b is that the number of students of each type matched to school b is proportional to $r_1 : \dots : r_{|T|}$. In that case, each school b has multiple levels over each type t and level number j , $\ell_{b,t}^j = [(j-1) * r_t, j * r_t - 1]$. For example, $\ell_{b,t}^1 = [0, r_t - 1]$ and $\ell_{b,t}^2 = [r_t, 2r_t - 1]$.

Egalitarian Egalitarian is a special case of proportionality where the ratio among all types is 1. In that case, each school b has multiple levels over each type t where each level $\ell_{b,t}^j = [j-1, j-1]$.

Our diversity framework can also capture combinations of the above objectives. For example, one could use minimum quotas to ensure that school b has a certain number of students from each type and beyond that we can use the egalitarian objective so that when all minimum quotas have been met, we then give the highest priority to the student whose type exceeds its minimum quota by the least margin.

Next, we give an example of how our approach can capture proportionality goals.

Example 1 (Proportional Diversity Goals). Consider a school b and a set of types $T = \{t_1, t_2, t_3, t_4\}$ with desired ratios 1 : 2 : 3 : 2. Then graphically R_b is depicted in Figure 2. Based on a current allocation of a school, each type has its own count of how many admitted students satisfy that type. Which types have the highest precedence for the school depends on current level of the types based on their counts. Suppose that the current allocation leads to the following counts $t_1 : 0, t_2 : 1, t_3 : 4, \text{ and } t_4 : 7$. Then, then precedence level of the types are as follows: $t_1 : 1, t_2 : 1, t_3 : 2, \text{ and } t_4 : 4$. Hence, t_1 and t_2 are the highest precedence types.




Figure 2: Proportional goals in Example 1.

4.2 Choice Function for Flexible Diversity Goals

For a given school b , its priority relation \succ_b over the students along with its precedence relation R_b over types can guide the school to select students from applicants. We specify a particular choice function for the schools.

We define a natural choice function of school b that chooses a set of contracts in Algorithm 1. The high level idea is that it identifies the types whose precedence is most important (based on which students are already selected) and then selects the highest priority contract that involves a student who satisfies one of the types. The process is repeated until the school capacity is reached or all students are selected.

Next, we explain Algorithm 1 in more detail. Given a set of contracts Y_b that pertain to school b , a priority ordering \succ_b over contracts and a precedence ordering R_b over types, the choice function Ch_b works as follows. In the beginning, we remove all unacceptable contracts from Y_b and initialize the set Z to be empty which is used to store the set of contracts selected by school b . Next, we repeat the following procedure until the set of contracts Y_b becomes empty or the school capacity is reached: First, we identify the set of types T' such that i) each type $t \in T'$ has the highest precedence based on $R_b(Z)$, and ii) there exists some contract $x = (i, h, b)$ from Y_b in which student i has some type from T' . Then, we scan over the set of contracts Y_b based on the priority ordering \succ_b of school b and select the first contract (i, h, b) such that student i has one of the highest precedence types. Finally, we update Z and Y accordingly and remove any contract y' from Y_b that is associated with either student i or seat h .

Note that our new choice function (Algorithm 1) is a generalization of the one by Ehlers *et al.* [2014], and it is designed for more general diversity goals instead of minimum and maximum quotas only. It is also well-defined for students having overlapping types. Algorithm 1 can also easily incorporate hard upper bounds for each type. All of our results go through with these hard upper bounds.

Input: A set of contracts Y_b, R_b , and \succ_b
Output: A set of contracts $Z \subseteq Y_b$

- 1: Remove unacceptable contracts from Y_b .
- 2: $Z \leftarrow \emptyset$ % Z stores the set of contracts chosen by school b .
- 3: **while** $Y_b \neq \emptyset$ and $|Z|$ is no more than school capacity **do**
- 4: Identify the set of types T' such that i) each type $t \in T'$ has the highest precedence based on $R_b(Z)$ and ii) there exists some contract $y = (i, h, b) \in Y_b$ with $T(i) \cap T' \neq \emptyset$.
- 5: Scan over Y_b based on priority ordering \succ_b and select the first contract $y = (i, h, b)$ such that $T(i) \cap T' \neq \emptyset$.
- 6: $Z \leftarrow Z \cup \{y\}, Y_b \leftarrow Y_b \setminus \{y\}$.
- 7: Remove all contracts from Y_b that involve i or h .
- 8: **return** Z

Algorithm 1: Choice function Ch_b of school b

5 Stable Matching under Flexible Diversity Goals

In this section, we propose a two-sided matching algorithm that deals with general diversity goals.

5.1 Generalized Deferred Acceptance with Flexible Diversity

The algorithm is well-defined whether students have strict preferences over all the school seats or whether they are indifferent between all the school seats of the same school. It is also well-defined even if students have multiple types.

Given a set of contracts X , let $Ch_i(X)$ denote the choice function of student i that selects her favorite contract among X_i . Let $Ch_b(X)$ denote the choice function of school b that selects a set of contracts among X_b , which is not necessarily unique. Let $Ch_S(X) = \bigcup_{i \in S} Ch_i(X)$ and $Ch_C(X) =$

$\bigcup_{b \in C} Ch_b(X)$ denote the choice functions of students S and schools C , respectively.

Armed with our specified choice function for schools (Algorithm 1), we consider the framework of the Generalized Deferred Acceptance (GDA) by Hatfield and Milgrom [2005] that works as follows. Each student first chooses her favorite contract from the set Y . Among all contracts proposed by students, each school then chooses a set of contracts. All contracts that are not selected by any school are removed from the set Y . Repeat these procedures until no more contract is removed.

Input: Instance I, Ch_S, Ch_C , a set of contracts Y
Output: An outcome $Z \subseteq Y$

```

1:  $Re \leftarrow \emptyset, X \leftarrow Y, Z \leftarrow \emptyset$ 
2: while  $X \neq Z$  do
3:    $X \leftarrow Ch_S(Y \setminus Re)$            % Students select contracts.
4:    $Z \leftarrow Ch_C(X)$                % Schools select contracts.
5:    $Re \leftarrow Re \cup (X \setminus Z)$    % Update rejected contracts.
6: return  $Z$ 

```

Algorithm 2: Generalized Deferred Acceptance (GDA)

We will refer to the GDA algorithm with the choice function defined in Algorithm 1 as *Generalized Deferred Acceptance with Flexible Diversity (GDA-FD)*.

Here is an example illustrating how GDA-FD works.

Example 2. *There are 4 students $S = \{1, \dots, 4\}$ where students $\{1, 2, 3\}$ belong to type t_1 and student 4 belongs to type t_2 . Suppose there is one school b that has 3 school seats $H = \{h_1, h_2, h_3\}$, a priority list $1, 2, 3, 4$ and precedence relation that requires proportionality ratio of 2: 1 between t_1 and t_2 . The preference profile of students are as follows: $h_1 \succ_1 h_2 \succ_1 h_3, h_2 \succ_2 h_3 \succ_2 h_1, h_3 \succ_3 h_2 \succ_3 h_1, h_3 \succ_4 h_1 \succ_4 h_2$. We abuse the notation of preferences and contracts, since there is only one school b . In the first round, students propose $(1, h_1), (2, h_2), (3, h_3), (4, h_3)$ and school b chooses $(1, h_1), (2, h_2), (4, h_3)$. In the second round, students propose $(1, h_1), (2, h_2), (3, h_2), (4, h_3)$ and school b chooses $(1, h_1), (2, h_2), (4, h_3)$. In the final round, students propose $(1, h_1), (2, h_2), (3, h_1), (4, h_3)$ and school b chooses $(1, h_1), (2, h_2)$, and $(4, h_3)$.*

Theorem 1. *The GDA-FD algorithm runs in polynomial-time and returns a feasible outcome that is non-wasteful, even if each student has multiple types.*

Next, we propose a stability concept for the setting of School Choice with Flexible Diversity Goals and Specialized Seats by taking the dynamic precedence of schools into account. The following Definition 1 captures a natural idea called *dynamic priority* proposed in a influential work on school choice with diversity goals [Ehlers *et al.*, 2014]. While our new definition is applicable to any market with flexible diversity goals. When diversity goals are not considered, it is equivalent to the standard stability concept.

Definition 1 (Stability). *Given a feasible outcome Y with $(i, h', b') \in Y$, a student i and a school b will form a blocking pair if $(i, h, b) \succ_i (i, h', b')$ and either i) the outcome $Y \cup \{(i, h, b)\} \setminus \{(i, h', b')\}$ is feasible; or ii) there exists*

a contract $(j, h, b) \in Y$ such that for all $t \in T(i)$ and all $t' \in T(j)$ and for the outcome $Y' = Y \setminus \{(j, h, b)\}$, one of the following conditions holds: a) $t \succ_b(Y') t'$; or b) $(i, h, b) \succ_b (j, h, b)$ and either $T(i) = T(j)$ or $t \succ_b(Y') t'$ holds. A feasible outcome is stable if there is no blocking pair.

We will try to understand how far stability can be achieved under various conditions.

5.2 No Specialized Seats

We first focus on the case in which each school has identical/homogenous school seats or students are indifferent between school seats of the same school. This model is still a significant generalization of the controlled school choice problem [Ehlers *et al.*, 2014] in which all the school seats are identical and schools impose soft minimum and maximum quotas on each type. In contrast, our model allows for much more general diversity goals including type-specific quotas. Our first result is the following Theorem 2.

Theorem 2. *When there are no specialized seats and each student has one type, GDA-FD is strategy-proof for the students and yields a stable outcome.*

Note that to run GDA-FD, students need to break ties lexicographically to derive a strict preference relation over contracts. The proof idea for Theorem 2 is as follows. We show that in our framework, even though the choice functions of schools capture complex diversity goals, they still satisfy two key properties called substitutability (SUB) and law of aggregate demand (LAD) as defined by Hatfield and Milgrom [2005]. Once these properties are established, we invoke a general result of Hatfield and Milgrom [2005] that the Generalized Deferred Acceptance is strategy-proof for the students and always yields a stable outcome.

Note that Theorem 2 does not hold when each student may have multiple types, as mentioned in Theorem 3.

Theorem 3. *When there are no specialized seats and each student has multiple types, GDA-FD is no longer strategy-proof for the students and does not always yield a stable outcome.*

6 Specialized Seats

In this section, we delve deeper into the case where students may distinguish between certain seats of a school. We study which positive properties of GDA-FD such as strategyproofness continue to hold under more general preferences, and present negative results that stable outcomes are not guaranteed to exist even under very restrictive conditions.

6.1 Case of a Global School

We first warm up by getting an in-depth understanding of GDA-FD for the case of a global school. The case of a global school is still an important setting that captures scenarios such as hiring of candidates by a company under diversity goals. We provide an alternative view of GDA-FD by showing that it is equivalent to a sequential allocation algorithm (Algorithm 3) under the assumption that there is a policy maker who treats the set of school seats as one global school b that has a global priority and a precedence ordering. Algorithm 3

can be viewed as a generalized version of serial dictatorship or sequential allocation in which dynamic diversity priorities are taken into account.

Input: An instance I with global priorities and precedence relation for school b
Output: An outcome $Z \subseteq \mathcal{X}$

- 1: **while** some student is unmatched and can get an unallocated school seat **do**
- 2: Identify the set of types T' with the highest precedence based on $R_b(Z)$ for which there exists some student who satisfies some type in T'
- 3: Among the students who satisfy some type in T' , find the student $i \in S$ who has the highest school priority based on \succ_b (which is achieved from \succsim_b by any fixed tie-breaking rule)
- 4: Assign student i her favorite contract $x = (i, h, b)$ among \mathcal{X} , $Z \leftarrow Z \cup \{x\}$
- 5: Remove the student i from the market, $S \leftarrow S \setminus \{i\}$
- 6: Remove the set of contracts involving student i and seat h , $\mathcal{X} \setminus (\mathcal{X}_i \cup \mathcal{X}_h)$
- 7: **return** Z

Algorithm 3: Sequential Allocation Under Dynamic Diversity Goals

An outcome X is *weakly Pareto optimal* if there is no other outcome X' such that all the students get a more preferred outcome. An algorithm is *type-strategyproof* if no student has an incentive to report a subset of her true types. Next, we summarize the properties of the Algorithm 3 in Theorem 4.

Theorem 4. *If there is one global school, even if students have multiple types, Algorithm 3 is strategyproof and type-strategyproof and the outcome returned by Algorithm 3 is weakly Pareto optimal.*

Next, we show Algorithm 3 has a strong connection with GDA-FD for the setting in which students have strict preferences over the school seats and there is exactly one school.

Theorem 5. *When there is exactly one school, the GDA-FD algorithm returns the same outcome as Algorithm 3 even if each student may have multiple types*

Proof. (Sketch) Assume Algorithm 2 terminates in m iterations. For each iteration $k \in [1, m]$, let Y^1, \dots, Y^m denote the set of contracts proposed by students. We prove that Algorithm 3 also returns the same outcome Y^m given the same input in the Appendix. \square

Although the GDA-FD remains strategyproof when there is one global school and each student has a strict preference over school seats, it does not return a stable outcome even if each student has one type, as mentioned in Theorem 6.

Theorem 6. *When each student has a strict preference over school seats, then the GDA-FD algorithm and Algorithm 3 do not return a stable outcome even if there is only one school and each student has one type.*

6.2 School-based Preferences

In this section, we consider the case that students have school-based preferences. A student has *school-based preferences* if all the school seats of one school are preferred over all the

Input: I, Ch_S, Ch_C , a set of contracts Y

Output: An outcome $Z \subseteq Y$

- 1: Let I' be the instance in which students are indifferent between school seats of the same school.
- 2: Let Y denote the outcome after applying GDA-FD to I' .
- 3: Let S_j denote the set of students who get matched to school b_j in the outcome Y .
- 4: **for** each $j \in [1, \dots, m]$ where $m = |C|$ **do**
- 5: Apply GDA-FD to the set of students S_j and school b_j to allocate the school seats of b_j to S_j , with respect to the original preferences \succ_S in I . Let the outcome be Y'_j .
- 6: **return** $Z = \bigcup_{j=1}^m Y'_j$.

Algorithm 4: Modified GDA-FD: Algorithm for school-based preferences

school seats of another school or vice versa. Note that students may have different preferences over school seats within the same school. For school-based preferences, we assume that all the school seats of a school are acceptable or none. The model of school choice with identical school seats is a special class of school-based preferences. The case of a single school is also a special case of school-based preferences.

Theorem 7. *For school-based preferences, GDA-FD is not strategy-proof for students even if each student has one type.*

Next, we design a new algorithm (Algorithm 4) for the case of school-based preferences which invokes GDA-FD twice. In the first stage of Algorithm 4, we first run GDA-FD on a modified instance I' in which school seats within a school are identical. This gives us information about which students are matched to which schools. In the second stage, we apply GDA-FD to each school and the set of students who are matched to that school based on students' true preferences over school seats. The second step determines which student receives which school seat. Algorithm 4 is strategyproof if each student has a single type as shown in Theorem 8.

Theorem 8. *For school-based preferences, when each student has one type, Algorithm 4 is strategy-proof for students.*

Next, we point out an impossibility result that the set of stable outcomes may be empty for school-based preferences even if under very restrictive conditions as shown in Theorem 9.

Theorem 9. *When each student has a strict preference over school seats, the set of stable outcomes may be empty even if each student has one type and all schools have the same priority ordering over students.*

7 Conclusions

We proposed a matching market model for residential markets with flexible diversity goals. We provided a clear understanding of under which conditions a stable matching is guaranteed to exist. Our diversity goal framework can be applied to many other settings in which priorities for different types dynamically change based on the current allocation. An interesting research direction is to identify other sufficient conditions for the guaranteed existence of stable outcomes under diversity goals.

References

- A. Abdulkadiroğlu and T. Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.
- S. Ahmadi, F. Ahmed, J. P. Dickerson, M. Fuge, and S. Khuller. An algorithm for multi-attribute diverse matching. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 3–9, 2020.
- H. Aziz, S. Gaspers, Z. Sun, and T. Walsh. From matching with diversity constraints to matching with regional quotas. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, pages 377–385, 2019.
- H. Aziz, S. Gaspers, and Z. Sun. Mechanism design for school choice with soft diversity constraints. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 153–159, 2020.
- S. Baswana, P. P. Chakrabarti, S. Chandran, Y. Kanoria, and U. Patange. Centralized Admissions for Engineering Colleges in India. In *Proceedings of the 20th ACM Conference on Economics and Computation*, pages 323–324, 2019.
- N. Benabbou, M. Chakraborty, X. Ho, J. Sliwinski, and Y. Zick. Diversity constraints in public housing allocation. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*, pages 973–981, 2018.
- J. Chen, R. Ganian, and T. Hamm. Stable matchings with diversity constraints: Affirmative action is beyond NP. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 146–152, 2020.
- L. Ehlers, I. E. Hafalir, M. B. Yenmez, and M. A. Yildirim. School choice with controlled choice constraints: Hard bounds versus soft bounds. *Journal of Economic Theory*, 153:648–683, 2014.
- Y. A. Gonczarowski, N. Nisan, L. Kovalio, and A. Romm. Matching for the Israeli “mechinot” gap year: Handling rich diversity requirements. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*, pages 321–321, 2019.
- I. E. Hafalir, M. B. Yenmez, and M. A. Yildirim. Effective affirmative action in school choice. *Theoretical Economics*, 8(2):325–363, 2013.
- J. W. Hatfield and P. R. Milgrom. Matching with contracts. *American Economic Review*, 95(4):913–935, 2005.
- Y. Kamada and F. Kojima. Efficient matching under distributional constraints: Theory and applications. *The American Economic Review*, 105(1):67–99, 2015.
- Y. Kamada and F. Kojima. Recent developments in matching with constraints. *The American Economic Review*, 107(5):200–204, 2017.
- F. Kojima. School choice: impossibilities for affirmative action. *Games and Economic Behavior*, 75(2):685–693, 2012.
- S. D. Kominers and T. Sönmez. Designing for diversity in matching. In *Proceedings of the 14th ACM conference on Electronic Commerce (EC)*, pages 603–604, 2013.
- S. D. Kominers and T. Sönmez. Matching with slot-specific priorities: Theory. *Theoretical Economics*, 11(2):683–710, 2016.
- R. Kurata, N. Hamada, A. Iwasaki, and M. Yokoo. Controlled school choice with soft bounds and overlapping types. *Journal of Artificial Intelligence Research*, 58:153–184, 2017.
- T. Nguyen and R. Vohra. Stable matching with proportionality constraints. *Operations Research*, 67(6):1503–1519, 2019.
- T. Sönmez and M. B. Yenmez. Constitutional implementation of reservation policies in India. Manuscript, 2019.
- Z. Sun. Matching with constraints. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), Doctoral Consortium*, pages 6464–6465, 2019.